Consistent Judgment Aggregation in Liquid Democracy: Utilizing Delegation Structure in the Ranked Agenda Rule

MSc Thesis *(Afstudeerscriptie)*

written by

Pelle ℕelissen
(born 20 October 1997 in Heemstede, the Netherlands)

under the supervision of **Prof. Dr. Ulle Endriss**, and submitted to the Examinations Board in partial fulfillment of the requirements for the degree of

MSc in Logic

at the *Universiteit van Amsterdam.*

Date of the public defense:	**Members of the Thesis Committee:**
26 March 2024	Dr. Benno van den Berg (chair)
	Dr. Davide Grossi
	Dr. Fernando P. Santos
	Prof. Dr. Ulle Endriss

INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

Acknowledgements

First and foremost, I would like to express my sincere gratitude to Ulle Endriss, my thesis supervisor. It was in the first month of the Master of Logic that he introduced me to the fascinating field of social choice theory. Throughout the program, he realistically guided me toward the socially relevant applications of the exact sciences. His dedicated, inspiring, and relaxed supervision of my master thesis has been invaluable. The Master of Logic would have been very different, had Ulle not shared his passion for this academic field with me on so many occasions.

Secondly, I extend massive thanks to the other members of my thesis committee, Benno van den Berg, Davide Grossi, and Fernando P. Santos, for their time spent reading my work, posing thought-provoking questions, and providing valuable feedback.

Lastly, I would like to thank Sabina and Gianluca. I wouldn't have wanted to experience this program without you both. Together, we studied numerous logics, attended multiple conferences, and embarked on obscure ferry trips to remote summer schools. Yet, what I'll treasure most are the beautiful conversations, the breathtaking mountains, and the delicious FEBO snacks we enjoyed. Graciis for everything!

Abstract

Liquid democracy is a collective decision-making process in which voters are allowed to delegate their vote to any peer. We consider a liquid version of judgment aggregation – a setting in which the collective decision concerns multiple interdependent issues. We translate a number of well-known normative requirements on aggregation mechanisms from the literature to our setting, and reproduce two famous impossibility theorems. One such mechanism is the ranked agenda rule, which efficiently computes a collective judgment by prioritizing issues that receive strong support from the voters. To arrive at satisfactory and logically consistent collective judgments in liquid democracy, we propose a refinement of the ranked agenda rule that breaks ties between issues with an equal number of supporters by the underlying delegation structure.

We motivate the rule formally by axiomatically characterizing the ranked agenda rule, and numerically by studying its behavior on a computational model of voter behavior. The computational model probabilistically simulates boundedly rational voters translating their uncertain preferences to liquid ballots. We study the correlation between delegation structure and epistemic performance of the profiles generated, and find that deeper delegation structures tend to approximate voters' true preferences less accurately. When we numerically compare our structural ranked agenda rule with Kemeny's rule, the original ranked agenda rule and viscous democracy, we conclude that all rules approximate the optimal collective decision equally accurately, but that our structural rule returns fewer possible collective judgments per input profile, i.e., is more resolute.

Contents

1 Introduction **4**
1.1 A Brief History of Liquid Democracy 5
1.2 In Defense of Liquid Democracy 5
1.3 Technical Challenges in Liquid Democracy 8
1.4 Objectives and Outline of the Thesis 10

2 Judgment Aggregation in Liquid Democracy **12**
2.1 A Model of Liquid Aggregation 12
2.2 Domain Restrictions and Normative Axioms 15
2.3 Three Impossibility Results 28

3 The Ranked Agenda Rule **32**
3.1 The Original Definition 32
3.2 Ranked Agenda as a Prudent Rule 36
3.3 A Prudent Characterization 41

4 Delegation Structure and Epistemic Performance **52**
4.1 Motivation 52
4.2 A Model of Profile Formation 53
4.3 Experiments and Results 61
4.4 Discussion 72

5 A Structural Ranked Agenda Rule **74**
5.1 Definition and Normative Properties 74
5.2 Profile Formation for Multiple Issues 77
5.3 Numerical Evaluation of Four Rules 79

6 Conclusion **85**

Bibliography **87**

Chapter 1

Introduction

"Democracy is both an ideal and a method ... [and] voting is at the heart of both the method and the ideal," Riker [1988] wrote in his *Liberalism Against Populism*. In the book, he points out that the theory of democracy and the theory of voting are inherently connected. Whereas most of democratic theory is developed by political scientists and economists, voting theory, and social choice theory more broadly, has become a significant and independent field of research at the intersection of the humanities and the exact sciences. Under the heading of Digital Democracy, an interdisciplinary community of political scientists, computer scientists, mathematicians, logicians and economists is exploring methods of democratic decision making that take advantage of the digital tools of the current century. One of the proposals is *liquid democracy*, a form of democracy that aims to combine the best of two worlds: the ideal of direct democracy and the practicality of representative democracy.

In liquid democracy, voters may choose either to cast a direct vote on the issue at hand, or to delegate their voting power to another voter. In the spirit of direct democracy, this allows each voter to directly express their opinion on the matter. Simultaneously, it releases voters of the responsibility to form an informed opinion on each and every topic: as in a representative democracy, a voter can always delegate their vote to a trusted peer.

This thesis studies the mechanisms (or *aggregation rules*) which convert a collection of votes in a liquid democracy instance to a collective decision. In particular, we propose an aggregation rule (the *structural ranked agenda rule*) which guarantees logically consistent collective decisions, while exploiting the delegation structure to break ties between different acceptable collective decisions. At the end of this chapter (Section 1.4), we elaborate on the objectives of our work and provide an outline of the thesis. But first, we present the background information and technical terms necessary to state our precise objectives by briefly introducing the history of liquid democracy (Section 1.1), presenting its most important merits (Section 1.2) and describing its major unresolved technical challenges (Section 1.3).

1.1 A Brief History of Liquid Democracy

Dodgson [1884] (alias Lewis Carroll) is generally considered to be the first to propose an idea similar to liquid democracy [Behrens, 2017]. He was primarily concerned with votes 'going to waste' if a representative in 19th century Britain received more votes than necessary to secure a seat in parliament. He proposed to allow the representative to transfer his votes to another parliamentary candidate of his own political party that received too few votes to enter parliament. This process was to be executed after all votes were counted and the result publicly announced.

It wasn't until eighty years later, that Dodgson's idea of delegating votes received new significant attention, when the invention of the modern computer opened up new possibilities for large-scale social deliberation [Behrens, 2017]. Tullock [1967] and Miller [1969] both proposed systems in which citizens could follow political debates via "broadcast" and could choose to delegate their vote to a representative, or to vote "by wire" or using a "console tied to a computer". Their proposals are very close to the concept understood as liquid democracy today, although they lack an important feature: the explicit authorization of delegates to re-delegate their received votes further, i.e., transitive delegation.

Ford [2002] was probably the first to explicitly allow for re-delegation of votes, when he proposed six basic axioms for his form of "delegative democracy" [Behrens, 2017]. Notably, Ford already proposed two extensions to his system, "backup votes" and "split delegation", which have later become important proposals to resolve the issue of cyclic delegation (i.e., delegations in which two or more voters delegate to each other in an infinite loop).

While the merits and drawbacks of multiple liquid democracy proposals are still actively being studied, the idea has already been applied in practice. Digital tools for 'interactive democracy' are currently being used for internal decision making in some political parties around the world and for online community engagement platforms [Behrens et al., 2014, Boella et al., 2018, De Cindio and Stortone, 2013, Mancini, 2015]. A famous example of such political parties are the Pirate Parties in 43 countries, that use liquid democracy for internal decision making [Litvinenko, 2012]. The region of Friesland in Germany is even experimenting with liquid democracy as a form of government, although Neuland [2014] reports that citizen participation is rather low.

1.2 In Defense of Liquid Democracy

As stated, it is promised that liquid democracy combines the best of direct democracy and representative democracy. Blum and Zuber [2016] attempt to justify this promise on the basis of normative democratic theory. They consider a form of liquid democracy that has at least the following four properties.

Direct democracy Voters are entitled to vote directly on all policy issues, without the intervention of a representative.

Flexible delegation Voters are entitled to delegate their vote to any other voter, both on an issue-by-issue basis, or for one or multiple policy areas at once.

Transitive delegation Voters are entitled to re-delegate the votes they received through delegation, to any other voter.

Instant recall Voters are entitled to terminate their delegation at any time.

Taking representative democracy as a normative benchmark, they argue that liquid democracy performs better from the perspective of epistemic democracy (the idea that democracy is a method to arrive at the *truth*) as well as from the perspective of procedural democracy (the idea that democracy derives it legitimacy from its egalitarian procedures).

1.2.1 An Epistemic and an Egalitarian Justification

The epistemic justification of liquid democracy relies on the premise that all voters have some inaccuracies in their perception of the common good, but some voters (the 'experts') hold more accurate beliefs than others *for the relevant policy area(s).* Since liquid democracy allows for policy area specific delegation (and representative democracy does not), liquid democracy has a "greater capacity for mobilizing policy area expertise" [Blum and Zuber, 2016]. Furthermore, delegations to representatives that prove unable to determine what policy serves the common good, can be recalled instantaneously. As a result, liquid democracy is able to "effectively and efficiently filter out the best experts without the time lags of electoral cycles" [Blum and Zuber, 2016].

In the egalitarian justification of liquid democracy, we assume that members of a political community vote on the basis of their subjective interests and a democratic system is *better* if it allows all voters to pursue their interests in a more equal fashion. Liquid democracy is better than representative democracy in this sense, firstly because its direct democracy property ensures that voters have the ability to directly vote for their subjective interests. Secondly, while representative democracy forces voters to choose a single representative for all policy areas, the flexible delegation of liquid democracy allows voters to express complex or uncommon opinions by combining policy bundles proposed by different policy experts. Consequently, voters are better able to match their delegations to their subjective interests. In other words, liquid democracy "remedies the problem of unequal participatory power that is generated in representative democracy by differences between how well the different members' interests match with available policy bundles" [Blum and Zuber, 2016].

1.2.2 An Information-Economic Formalization

In a similar attempt to justify liquid democracy, Green-Armytage [2015] provides an axiomatic and an information-economic argument. Axiomatically, he concludes that a system of "voluntary delegation" is necessary on the basis of the axioms that (1) in a democracy, voters should have the right to vote directly on any issue, and (2) voters have the right to delegate their vote to any peer. The former axiom is assumed to be a "familiar notion", while the latter is justified as saving time compared to informing oneself and going out to vote, and as being equivalent to the free flow of information: if we view a delegation not as a flow of power to the representative, but as a "transfer of information (about how to vote)" in the opposite direction, then the rejection of the latter axiom, is a restriction on the flow of information.

His information-economic justification effectively formalizes the arguments of Blum and Zuber [2016]. It relies on a spacial imperfect information model: each possible position on an issue is represented by a real number and all voters occupy some point in a multi-dimensional Euclidean space, where each dimension represents some issue. Furthermore, each voter's uncertainty about their own position (i.e., their perfect self-interested direct vote) and about the position of all other voters are modeled as random variables. Based on this imperfect information, voters are tasked with casting a vote that represents their true position as accurately as possible. We compare their ability to do so between direct, liquid and representative democracy.

In direct democracy, voters have no choice but to vote according to their perceived subjective position. In representative democracy, voters are forced to select a single representative that minimizes the perceived Euclidean distance between the voter and the representative in this multi-dimensional space. In liquid democracy, voters may delegate their vote on any single issue to a representative, if the perceived one-dimensional distance on that issue between voter and representative, is smaller than the expected perception-error in the voter's own position on the issue. Green-Armytage [2015] mathematically shows that this option increases voters' ability to accurately reflect their true opinion, compared to direct democracy as well as to representative democracy. In other words, giving voters the option to delegate to different representatives on different issues, allows them to pursue their self-interest more accurately than in direct or representative democracy. This formalizes the egalitarian argument of Blum and Zuber [2016]. Furthermore, the median of all votes in the liquid democracy setting is closer to the median opinion of all voters, than in the direct and representative democracy settings. This in turn, formalizes the epistemic argument of Blum and Zuber [2016].

The fact that liquid democracy performs better than direct democracy in the analysis of Green-Armytage [2015], might surprise the reader with a background in social choice theory, but follows from the assumption that voters are uncertain about their own self-interest. Therefore, it might be the case that a delegate approximates the voter's self-interest better than the voter does herself. For example, Alice might know very little about climate policy and thus does not know which policy option reflects her values, while she does share many of here core values with Bob, who is a climate scientist. Delegating her vote on climate policies to Bob, might therefore serve Alice's self-interest better than voting directly.

1.2.3 Overburdening the Citizen

A natural concern is that liquid democracy (like direct democracy) overburdens the average citizen. Blum and Zuber [2016] point out that the "informational demands placed on the individual in a liquid democracy appear to be much greater than those placed on a voter in a representative democracy": voters must decide for each individual issue or policy area whether to vote directly (and if so, how to vote) or to delegate (and if so, to whom). The question arises if the average citizen is capable of performing these tasks effectively. Social epidemiologists argue that laypersons are well capable of evaluating the performance of experts without becoming an expert themselves [Blum and Zuber, 2016]. An important tool for doing so are "reputation systems" that rank or rate experts and are developed by an epistemic community. And even if expert selection fails, liquid democracy – more than representative democracy – benefits from the 'wisdom of crowds': the theory that the wisdom of many non-experts can jointly be more accurate than the judgment of any single expert. Finally, there is empirical evidence that more possibilities for political participation causes citizens to become more informed on political issues [Blum and Zuber, 2016]. Therefore, as citizens get used to a system of liquid democracy, they are likely to become better at casting the 'right' vote or selecting the 'right' delegate.

1.2.4 Conclusion

In conclusion, the core idea of liquid democracy, voluntary and flexible delegation, can be justified on the basis of procedural as well as epistemic democratic theory. Moreover, the political philosophical argument can be supplemented by an axiomatic argument and an information-economic formalization. Finally, there is reason to believe that the average citizen is not overburdened by her responsibilities in liquid democracy. Nonetheless, multiple conceptual, technical and implementation challenges remain in the theory of liquid democracy.

1.3 Technical Challenges in Liquid Democracy

The literature explores many major and minor problems with liquid democracy, both of political, conceptual or philosophical nature, and of technical nature. We focus on the main technical issues in the realm of mathematics, computer science and logic. We refer to the work of Brill [2018] for a more substantive overview of the technical challenges faced by liquid democracy, and to Blum and Zuber [2016] for some politically philosophical challenges.

1.3.1 Cyclic Delegation and Abstention

Cyclic delegation refers to the case where two or more voters delegate to each other in a cycle. Clearly, if this happens, we cannot count their votes towards a social decision without some further information on their preferences. A similar situation arises when a voter delegates to a peer, who abstains from voting for whatever reason. Christoff and Grossi [2017] show that (under the assumption that any delegation profile is equally likely), the probability that *all* voters (indirectly) delegate to a cycle, goes to $\frac{1}{e^2} \approx 14\%$ as the number of voters goes to infinity. Although this underlying assumption is rather unrealistic, we clearly need some method of resolving delegation cycles (and abstentions).

A straightforward way of doing so, is asking all voters to submit a 'backup vote', which is used in case their representative abstains or is part of a delegation cycle. However, this requires all voters to submit ballots on every individual issue (and thus prohibits automatic delegation on all issues within some policy area(s)), thereby losing many of the practical advantages of liquid democracy. Another common proposal is asking all voters to submit 'backup delegations'. How these additional delegations are used in the decision procedure differs per proposal.

Kotsialou and Riley [2020], Brill et al. [2022] and Colley et al. [2022] suggest multiple approaches where some search algorithm (most intuitively, breadth first search) decides which delegation option is used for each voter. The final delegations are fixed one by one, and if a preferred delegation of some voter leads to a cycle, her next alternative delegation is used instead.

Brill [2018] and Utke and Schmidt-Kraepelin [2023] suggest considering the delegation graph as a Markov chain, where primary delegations have a weight proportional to ϵ, secondary delegations to ϵ^2, etc. Utke and Schmidt-Kraepelin [2023] show that this approach allows for aggregation rules that are *copy-robust*, the property that voters cannot manipulate the outcome by copying their preferred delegate's vote instead of submitting a ranked delegation.

Other approaches of preventing or resolving cyclic delegation include the suggestion of Kahng et al. [2021] to assign voters a competency level and only allowing delegations to more competent peers, and the suggestion of Escoffier et al. [2019] to allow voters to change their vote iteratively in case they delegated to a cycle.

1.3.2 Policy-Inconsistency

It is well-known in the field of judgment aggregation, that even if all voters hold consistent beliefs, collective judgments formed using reasonable aggregation rules, may be inconsistent. The most famous example is probably the *doctrinal paradox* introduced by Kornhauser and Sager [1993]. Judgment aggregation theory teaches us that such collective inconsistencies can occur in almost any reasonable system containing a direct democracy component (see, e.g., the survey by List [2012]). Representative democracy, on the other hand, "is strong at bundling together solutions to societal problems from different policy areas—

and also trading them against each other, if necessary," due to politicians (and their parties) being involved in multiple policy areas simultaneously, Blum and Zuber [2016] argue. If liquid democracy is to be a viable alternative to representative democracy, the issue of policy-inconsistency must thus be addressed. Even if the possibility of collective inconsistency cannot be ruled out for *any* combination of votes, limiting inconsistencies as much as possible, is an important challenge.

Christoff and Grossi [2017] treat delegation graphs in liquid democracy as a social network and model the stabilization of opinions, if voters copy the opinion of their delegate iteratively. They enforce that voters only copy an opinion if it is consistent with their opinions on other issues. As a result, the stable opinions are all individually consistent. Applying an aggregation rule from judgment aggregation theory that preserves consistency, will thus yield a consistent collective opinion. However (under some technical assumptions on the issue-agenda, and independence and unanimity), the only such rules are oligarchies, in which a subset of the voters have full authority over the collective decision. This is a well-known result in judgment aggregation theory by Dietrich and List [2008], that Christoff and Grossi [2017] generalized to the liquid democracy setting.

In the special case of ordinal voting (i.e., the setting in which voters rank a number of alternatives and one alternative is elected), Brill and Talmon [2018] show that the detection of individual inconsistencies in a delegation graph is NP-complete. To recover individual consistency, they suggest using distance rationalization on delegation structures, in the spirit of Kemeny's rule that minimizes pairwise swaps of candidates in individual ballots to obtain unanimity [Kemeny, 1959]. For a given delegation graph, we search for the closest consensus graph (a delegation graph with a 'clear winner') according to some distance metric between delegation graphs. The winner of this consensus graph is declared the winner of the vote.

1.3.3 Other Challenges

Many other challenges in liquid democracy are described in the literature. For example, Green-Armytage [2015] and Blum and Zuber [2016] make a number of suggestions for practical implication, including the role of the executive branch in relation to a liquid legislature. Brill [2018] and Harding [2022] address some aspects of voter manipulation and strategyproofness. And finally, Harding [2022] considers questions of monotonicity, i.e., whether increased support can harm an alternative.

1.4 Objectives and Outline of the Thesis

In this thesis, we will focus our attention on the issue of policy-inconsistency. The objective is to construct an aggregation rule for liquid democracy that guarantees an outcome that is consistent, while meeting as many desirable fairness constraints as possible. In our setting, voters will simultaneously express their opinion on multiple issues, which are logically connected through some ‘integrity constraint’. In order to keep our formal results as general as possible, we do not assume that voters’ ballots are in any way consistent with the integrity constraint or that delegations are always acyclic, but we do assume that voters cannot abstain from voting.

However, when assessing the effectiveness of aggregation rules, we strengthen our assumptions. That is, we will assume that voters respect the integrity constraint in the *direct* votes they submit, while allowing them to delegate their votes on remaining issues freely: we will not assume, that the opinion of a voter’s delegate on one issue is consistent with the voter’s direct vote on another issue. We believe that these assumptions mimic the reality of liquid democracy as closely as possible: we may generally assume that voters are rational (and thus that their direct votes on several issues are logically consistent), but we may not assume that they can predict the exact vote of their delegates (and thus their delegates’ votes might be inconsistent with their own direct votes).

The thesis is structured as follows. In Chapter 2, we define our formal model of judgment aggregation in liquid democracy and its normative axioms, and generalize some basic results from judgment aggregation theory. In Chapter 3, we define our main aggregation mechanism, the ranked agenda rule, and characterize it in terms of normative axioms. In Chapter 4, we define a computational model of voter behavior in liquid democracy, and use it to study the connection between the delegation structure of a liquid democracy instance and its epistemic performance. Finally, in Chapter 5, we combine the results from the previous chapters to obtain a structural ranked agenda rule, and computationally compare its epistemic performance against three other aggregation rules.

Chapter 2

Judgment Aggregation in Liquid Democracy

We begin this chapter by defining our model of liquid democracy in Section 2.1. In Section 2.2, we introduce a number of domain restrictions and normative axioms. Finally, we translate some basic impossibility results from judgment aggregation theory to our setting in Section 2.3.

2.1 A Model of Liquid Aggregation

Following the assumptions in Section 1.4, we define a model of judgment aggregation in liquid democracy as follows. Any liquid democracy instance consists of a finite set $\mathcal{N} = \{1, \ldots, n\}$ of *voters* and a finite set $\Phi^+ = \{p_1, \ldots, p_m\}$ of *issues*, modeled as propositional letters. The set $\Phi^- = \{\neg p_1, \ldots, \neg p_m\}$ contains all negations of the propositional letters and we call $\Phi := \Phi^+ \cup \Phi^-$ the *agenda* of the liquid democracy instance. Elements of the agenda are referred to as *literals*. For any literal $\ell \in \Phi$, we write $\sim\ell$ for $\neg p$ if $\ell = p \in \Phi^+$, or for p if $\ell = \neg p \in \Phi^-$. Furthermore, let $\mathcal{L}_{\text{sat}}$ be the set of satisfiable propositional formulas over the propositional letters in Φ^+. This set contains every possible *integrity constraint* $\Gamma \in \mathcal{L}_{\text{sat}}$, expressing the logical connections between issues.

Definition 2.1.1 (Judgment Set)**.** We call a set $J \subseteq \Phi$ a *judgment set*. We call a judgment set *complete* if it contains p or $\neg p$ for all $p \in \Phi^+$, and we call it *complement-free* if it does not contain both p and $\neg p$ for any $p \in \Phi^+$. For $\Gamma \in \mathcal{L}_{\text{sat}}$, we say that a judgment set J is *consistent with* Γ, if $J \cup \{\Gamma\}$ is logically consistent. The set of all judgment sets that are complete and consistent with Γ, is denoted by $\mathcal{J}(\Gamma)$.

A judgment set represents an opinion that a voter (or a collective of voters) might have on the issues of the liquid democracy instance. If a judgment set is complete, it expresses a judgment on each issue, and if it is complement-free, it does not express contradictory judgments. If $\Gamma \in \mathcal{L}_{\text{sat}}$ is an integrity constraint, a judgment set that is consistent with Γ, expresses an opinion that respects the logical connections between the issues which are modeled by Γ.

In a liquid democracy instance, all voters submit a *ballot* stating their direct votes on some issues and their delegations to other voters on the remaining issues. Ballots are formally modeled by the following definition.

Definition 2.1.2 (Liquid Aggregation Ballot)**.** A *liquid aggregation ballot* (or *ballot* for short) for voter $i \in \mathcal{N}$ is a map

$$B_i \colon \Phi \to \{+, -\} \cup (\mathcal{N} \setminus \{i\}),$$

such that for all $\ell \in \Phi$, we have $B_i(\ell) = +$ if and only if $B_i(\sim\ell) = -$, and if $B_i(\ell) \in \mathcal{N} \setminus \{i\}$, then $B_i(\ell) = B_i(\sim\ell)$. The set of all possible ballots for voter $i \in \mathcal{N}$ is denoted by $\mathcal{B}_i^\Phi$.

If $\ell \in \Phi$ is a literal and $B_i \in \mathcal{B}_i^\Phi$ is a ballot for voter $i \in \mathcal{N}$, then $B_i(\ell) = +$ (or $B_i(\ell) = -$) represents a direct vote in favor of (or against) literal ℓ (which automatically implies a direct vote against (or in favor of) $\sim\ell$); and $B_i(\ell) = j$ for $j \in \mathcal{N} \setminus \{i\}$ represents the delegation of i's vote on ℓ to another voter j (which automatically implies the delegation of i's vote on $\sim\ell$ to voter j). Thus, the set $B_i^{-1}(+)$ contains exactly those literals, that the voter accepts by a direct vote. Note that this is a complement-free judgment set (but not necessarily a complete one, or one that is consistent with the integrity constraint).

The combination of all voters' ballots is referred to as a *profile*, which is defined as follows.

Definition 2.1.3 (Liquid Aggregation Profile)**.** A *liquid aggregation profile* (or *profile* for short) for voters $\mathcal{N}$ is a vector $\boldsymbol{B} = (B_1, \dots, B_n)$ containing a ballot $B_i \in \mathcal{B}_i^\Phi$ for each voter $i \in \mathcal{N}$. The set of all possible profiles for voters $\mathcal{N}$ is denoted by $\mathcal{B}_\mathcal{N}^\Phi$. The set of all possible profiles for any non-empty finite set of voters is denoted by $\mathcal{B}^\Phi$.

In some parts of this thesis, we will need profiles consisting of different numbers of voters. If it is not clear from the context which set of voters we are considering, we denote by $\mathcal{N}_{\boldsymbol{B}}$ the set of voters that constitute the profile $\boldsymbol{B} \in \mathcal{B}^\Phi$. However, in most cases, all profiles can be considered to have some generic and fixed set of voters $\mathcal{N}$. Furthermore, if we view the delegations concerning literal $\ell \in \Phi$ in a profile as a graph, then for each voter $i \in \mathcal{N}$ there is a unique path from that voter to a voter who votes directly on ℓ, or to a delegation cycle (i.e., a subset of voters, who delegate to each other in a loop). We write $\text{path}_{\boldsymbol{B},\ell}(i)$ for the set containing the voters in this path, and $\text{rep}_{\boldsymbol{B},\ell}(i)$ for voter i's representative for ℓ (if she exists), i.e., the unique voter in $\text{path}_{\boldsymbol{B},\ell}(i)$ that casts a direct vote on ℓ (if she exists). In particular, if i votes directly on ℓ, then $\text{path}_{\boldsymbol{B},\ell}(i) = \{i\}$ and $\text{rep}_{\boldsymbol{B},\ell}(i) = i$. If $\text{rep}_{\boldsymbol{B},\ell}(i)$ exists for all $i \in \mathcal{N}$ and all $\ell \in \Phi$, then we call profile $\boldsymbol{B}$ *acyclic*.

The mechanism that aggregates the voters' ballots into a single collective decision, is referred to as a *liquid aggregation rule* and is formally defined as follows.

Definition 2.1.4 (Liquid Aggregation Rule)**.** A *liquid aggregation rule* $\mathcal{F}$ (or *rule* for short) with *domain* $\mathcal{D}_{\mathcal{F}} \subseteq \mathcal{B}^{\Phi} \times \mathcal{L}_{\text{sat}}$ is a map $\mathcal{F} \colon \mathcal{D}_{\mathcal{F}} \to \mathcal{P}(\mathcal{P}(\Phi)) \setminus \{\emptyset\}$.

For rule $\mathcal{F}$, profile $\boldsymbol{B} \in \mathcal{B}^{\Phi}$ and integrity constraint $\Gamma \in \mathcal{L}_{\text{sat}}$ such that $(\boldsymbol{B}, \Gamma) \in \mathcal{D}_{\mathcal{F}}$, a judgment set $J \in \mathcal{F}(\boldsymbol{B}, \Gamma)$ is said to be a *possible collective decision*. If $\mathcal{F}(\boldsymbol{B}, \Gamma) = \{J\}$, then we call J the *collective decision*.

We say that a literal $\ell \in \Phi$ is *necessarily accepted* (or *necessarily rejected*) if $\ell \in J$ (or $\ell \notin J$) for all $J \in \mathcal{F}(\boldsymbol{B}, \Gamma)$, and *possibly accepted* (or *possibly rejected*) if $\ell \in J$ (or $\ell \notin J$) for some $J \in \mathcal{F}(\boldsymbol{B}, \Gamma)$.

A liquid aggregation rule is a mechanism, which takes as input a profile and an integrity constraint, and returns a non-empty set of 'acceptable' collective decisions.[1] Thus, a liquid aggregation rule need not be *resolute*: it may return several possible collective decisions for a single input. Furthermore, since we define the integrity constraint to be an explicit input of the rule, a single liquid aggregation rule represents, in effect, a whole family of rules, one for each integrity constraint. In practice, one would expect the rule to behave similarly for similar integrity constraints, which can be formalized by normative axioms that constrain the outcome of liquid aggregation rules under varying integrity constraints (see Section 2.2).

Also note that a liquid aggregation rule comes with a specific domain on which it is defined. This allows us to consider different domain restrictions, which we define in Section 2.2.1. Finally, liquid aggregation rules are defined for all possible non-empty finite sets of voters simultaneously (i.e., not only for a fixed set $\mathcal{N}$), which is necessary for the characterization result in Section 3.3.

The field of judgment aggregation has developed a wide range of aggregation rules for the setting of direct democracy. Those rules can be generalized to liquid democracy rules in a standard way, if we convert the liquid profile to a profile containing only direct votes. We do so by following the delegation chains. That is, if $\boldsymbol{B} = (B_1, \ldots, B_n) \in \mathcal{B}^{\Phi}$ is a liquid profile, we call $\boldsymbol{B}' = (B'_1, \ldots, B'_n) \in \mathcal{B}^{\Phi}_{\mathcal{N}_{\boldsymbol{B}}}$ the *proxy profile of* $\boldsymbol{B}$ if for each voter $i \in \mathcal{N}_{\boldsymbol{B}}$ and literal $\ell \in \Phi$, we have $B'_i(\ell) = B_{\text{rep}_{\boldsymbol{B},\ell}(i)}(\ell)$. In case some $\text{rep}_{\boldsymbol{B},\ell}(i)$ is undefined (i.e., in case the delegations on some issue are cyclic), then no proxy profile exists.[2] After constructing a proxy profile (if it exists), any classical judgment aggregation rule can be applied to obtain a social decision, thus generalizing the classical judgment aggregation rule to a liquid aggregation rule. And the other way around, liquid aggregation problems can be reduced to classical judgment aggregation problems by considering the proxy profile.

[1]Throughout this thesis, we use the terms 'collective decision', 'collective judgment' and 'collective opinion' interchangeably.

[2]As discussed in Section 1.3.1, several solutions to cyclic delegation have been proposed in the literature. For simplicity, we choose to leave the proxy profile undefined in case of cyclic delegation.

2.2 Domain Restrictions and Normative Axioms

The model defined above is a generalization of the classical binary judgment aggregation setting (see, e.g., the survey by List [2012]): if all voters submit direct votes only, we are in a situation of direct democracy. And binary judgment aggregation is, in turn, a generalization of ordinal preference aggregation [Dietrich and List, 2007]. Social choice theorists have developed many normative axioms for aggregation rules in these (and other) settings. We will translate these axioms to our setting, and introduce new axioms capturing the behavior of liquid aggregation rules specifically. But we first consider some domain restrictions.

2.2.1 Domain Restrictions

In some cases, it is interesting to consider a liquid aggregation rule that is defined on a restricted domain. For example, some rules might always return consistent collective decisions on some domains, but not on others. Thus, we define a number of useful domain restrictions in this section. Note that a domain specifies the pairs of profiles and integrity constraints that a rule should (at least) be defined on. Thus, a rule might be defined for a profile $\boldsymbol{B}$ in combination with an integrity constraint Γ, but not for the same profile $\boldsymbol{B}$ in combination with a different integrity constraint Γ'; and the other way around.

If we allow voters to submit any ballot as defined in Definition 2.1.2, we are in the universal domain.

Definition 2.2.1 (Universal Domain)**.** A liquid aggregation rule $\mathcal{F}$ has the *universal domain*, if $\mathcal{D}_{\mathcal{F}}$ is the whole set $\mathcal{B}^{\Phi} \times \mathcal{L}_{\text{sat}}$.

As we argued in Section 1.4, it is interesting to assume a weak form of individual rationality, namely that the direct votes of each individual voter are consistent with the integrity constraint. This assumption restricts the domain as follows.

Definition 2.2.2 (Rational Domain)**.** A liquid aggregation rule $\mathcal{F}$ has the *rational domain*, if $\mathcal{D}_{\mathcal{F}}$ contains all pairs $(\boldsymbol{B}, \Gamma) \in \mathcal{B}^{\Phi} \times \mathcal{L}_{\text{sat}}$ such that for each voter $i \in \mathcal{N}_{\boldsymbol{B}}$, the set $B_i^{-1}(+)$ is consistent with Γ.

When we consider all acyclic profiles as valid inputs to a rule, we are in the acyclic domain, which is defined as follows.

Definition 2.2.3 (Acyclic Domain)**.** A liquid aggregation rule $\mathcal{F}$ has the *acyclic domain*, if $\mathcal{D}_{\mathcal{F}}$ contains all pairs $(\boldsymbol{B}, \Gamma) \in \mathcal{B}^{\Phi} \times \mathcal{L}_{\text{sat}}$ such that $\boldsymbol{B}$ is acyclic.

Finally, we consider the domain restriction that defines direct democracy. Note that the direct domain is a subset of the acyclic domain.

Definition 2.2.4 (Direct Domain)**.** A liquid aggregation rule $\mathcal{F}$ has the *direct domain*, if $\mathcal{D}_{\mathcal{F}}$ contains all pairs $(\boldsymbol{B}, \Gamma) \in \mathcal{B}^{\Phi} \times \mathcal{L}_{\text{sat}}$ such that $B_i(\ell) \in \{+, -\}$ for all literals $\ell \in \Phi$ and all voters $i \in \mathcal{N}_{\boldsymbol{B}}$.

Throughout this thesis, we will sometimes consider domains which are an intersection of the domains defined here. We will then concatenate the names of the domains. For example, we say that a rule $\mathcal{F}$ has the acyclic rational domain, if $\mathcal{D}_{\mathcal{F}}$ contains all pairs $(\boldsymbol{B}, \Gamma) \in \mathcal{B}^{\Phi} \times \mathcal{L}_{\text{sat}}$ such that $\boldsymbol{B}$ is acyclic and the direct votes of each voter are consistent with Γ.

Note that all domain restrictions allow the domain of a rule to be larger than strictly necessary to satisfy the definitions given. For example, any rule $\mathcal{F}$ that has the acyclic domain also has the direct domain, since the direct domain is a subset of the acyclic domain. Therefore, possibility results (i.e., theorems stating that a rule with certain properties exists) are stronger when shown on larger domains (optimally the universal domain), and impossibility results (i.e., theorems stating that no rules with certain properties exist) are stronger when shown on smaller domains.

2.2.2 Classical Normative Axioms

We turn our attention to the axioms of liquid democracy. We first consider some axioms which are primarily normative and can be used to describe the more common desirable properties of liquid aggregation rules. Most of these normative axioms are translations of famous axioms from judgment aggregation theory (see, e.g., the overviews by List [2012] and Endriss [2016]). Thereafter, in Section 2.2.3, we define the axioms that can be used to characterize the ranked agenda rule, which we study in Chapter 3. Although the latter axioms still express normative properties, they are more technical in nature.

Our first axiom distinguishes rules which always return a single collective decision (i.e., resolute rules), from rules which return several possible collective decisions in some cases (i.e., irresolute rules).

Axiom 2.2.5 (Resoluteness)**.** We call a liquid aggregation rule $\mathcal{F}$ *resolute*, if for any pair $(\boldsymbol{B}, \Gamma) \in \mathcal{D}_{\mathcal{F}}$, we have $|\mathcal{F}(\boldsymbol{B}, \Gamma)| = 1$. We call all other rules *irresolute*.

Unlike in our setting, aggregation rules in classical judgment aggregation theory are sometimes considered to be defined for a single integrity constraint. That is, if we change the integrity constraint, we move to a different aggregation rule. In order to generalize results from that setting, we must be able to restrict our attention to those liquid aggregation rules that are independent of integrity constraints. We call such aggregation rules *static*. Well-known examples of static rules are quota rules and dictatorships, whereas Kemeny's rule and the ranked agenda rule (see Section 3.1) are not static.

Axiom 2.2.6 (Staticity)**.** We call a liquid aggregation rule $\mathcal{F}$ *static*, if for all pairs $(\boldsymbol{B}, \Gamma_1), (\boldsymbol{B}, \Gamma_2) \in \mathcal{D}_{\mathcal{F}}$, we have $\mathcal{F}(\boldsymbol{B}, \Gamma_1) = \mathcal{F}(\boldsymbol{B}, \Gamma_2)$.

Another consequence of the introduction of the integrity constraint as an argument of aggregation rules is that it technically allows us to use different aggregation methods, even if two integrity constraints are logically equivalent. The axiom of language independence prohibits this peculiarity.

Axiom 2.2.7 (Language Independence)**.** We call a liquid aggregation rule $\mathcal{F}$ *language-independent*, if for all pairs $(\boldsymbol{B}, \Gamma_1), (\boldsymbol{B}, \Gamma_2) \in \mathcal{D}_{\mathcal{F}}$, we have $\mathcal{F}(\boldsymbol{B}, \Gamma_1) = \mathcal{F}(\boldsymbol{B}, \Gamma_2)$, whenever Γ_1 and Γ_2 are logically equivalent.

The following axioms require collective decisions to be complete, complement-free or consistent with the integrity constraint.

Axiom 2.2.8 (Collective Completeness)**.** We call a liquid aggregation rule $\mathcal{F}$ *(collectively) complete*, if for any pair $(\boldsymbol{B}, \Gamma) \in \mathcal{D}_{\mathcal{F}}$, all possible collective decisions $J \in \mathcal{F}(\boldsymbol{B}, \Gamma)$ are complete.

Axiom 2.2.9 (Collective Complement-Freeness)**.** We call a liquid aggregation rule $\mathcal{F}$ *(collectively) complement-free*, if for any pair $(\boldsymbol{B}, \Gamma) \in \mathcal{D}_{\mathcal{F}}$, all possible collective decisions $J \in \mathcal{F}(\boldsymbol{B}, \Gamma)$ are complement-free.

Axiom 2.2.10 (Collective Consistency)**.** We call a liquid aggregation rule $\mathcal{F}$ *(collectively) consistent*, if for any pair $(\boldsymbol{B}, \Gamma) \in \mathcal{D}_{\mathcal{F}}$, all possible collective decisions $J \in \mathcal{F}(\boldsymbol{B}, \Gamma)$ are consistent with Γ.

For the next axiom, we need some additional notation. Let $\sigma \colon \mathcal{N} \to \mathcal{N}$ be a permutation and let $\mathrm{id} \colon \{+, -\} \to \{+, -\}$ be the identity map. We denote by $\sigma \cup \mathrm{id}$ the map $\mathcal{N} \cup \{+, -\} \to \mathcal{N} \cup \{+, -\}$ that acts as σ on $\mathcal{N}$ and as id on $\{+, -\}$. If $\boldsymbol{B} = (B_1, \ldots, B_n)$ is a profile, we denote by $\sigma(\boldsymbol{B})$ the permuted profile $((\sigma \cup \mathrm{id}) \circ B_{\sigma^{-1}(1)}, \ldots, (\sigma \cup \mathrm{id}) \circ B_{\sigma^{-1}(n)})$. In effect, $\sigma(\boldsymbol{B})$ is the same profile as $\boldsymbol{B}$, but each voter $\sigma(i)$ plays the role in $\sigma(\boldsymbol{B})$ that voter $i \in \mathcal{N}$ played in $\boldsymbol{B}$. That is, we effectively shuffle the name tags of all voters. The axiom of anonymity states that such permutations should not change the outcome of a vote.

Axiom 2.2.11 (Anonymity)**.** We call a liquid aggregation rule $\mathcal{F}$ *anonymous*, if for any pair $(\boldsymbol{B}, \Gamma) \in \mathcal{D}_{\mathcal{F}}$ and any permutation $\sigma \colon \mathcal{N} \to \mathcal{N}$, we have $\mathcal{F}(\boldsymbol{B}, \Gamma) = \mathcal{F}(\sigma(\boldsymbol{B}), \Gamma)$.

A similar axiom exists for shuffling the name tags of propositional letters: the axiom of neutrality. Formally, the axiom states that if two propositional letters are treated identically by all voters, then any possible collective decision should either accept both propositional letters, or neither (and analogously for the negations of the propositional letters).

Axiom 2.2.12 (Neutrality)**.** We call a liquid aggregation rule $\mathcal{F}$ *neutral*, if for any pair $(\boldsymbol{B}, \Gamma) \in \mathcal{D}_{\mathcal{F}}$, any propositional letters $p, q \in \Phi^+$ and any possible collective decision $J \in \mathcal{F}(\boldsymbol{B}, \Gamma)$, we have that $p \in J$ if and only if $q \in J$, and $\neg p \in J$ if and only if $\neg q \in J$, whenever $B_i(p) = B_i(q)$ for all voters $i \in \mathcal{N}_{\boldsymbol{B}}$.

This is (an irresolute version of) the most common notion of neutrality: if we shuffle the names of propositional letters, this should not change the outcome of an aggregation rule. However, if we shuffle the names of the propositional letters in an aggregation profile, we must shuffle their names in the integrity constraint accordingly, if we want to preserve their roles in the aggregation problem as a whole. In other words, we should take into account the logical context of an aggregation instance, when considering neutrality. Therefore, we propose an axiom of *contextual neutrality*, which uses the following notation.

If $p, q \in \Phi^+$ and $\Gamma \in \mathcal{L}_{\text{sat}}$, we denote by $\Gamma[p \leftrightarrow q]$ a variant of formula Γ where all occurrences of p are replaced by q, and all occurrences of q by p (and thus all occurrences of $\neg p$ by $\neg q$, and $\neg q$ by $\neg p$). Similarly, if $J \subseteq \Phi$ is a judgment set, we denote by $J[p \leftrightarrow q]$ a variant of J where all occurrences of p are replaced by q, and all occurrences of q by p (and thus all occurrences of $\neg p$ by $\neg q$, and $\neg q$ by $\neg p$).

Note that our formulation of the axiom of neutrality is equivalent to requiring that for any possible collective decision $J \in \mathcal{F}(\boldsymbol{B}, \Gamma)$, we have that $J = J[p \leftrightarrow q]$, whenever $B_i(p) = B_i(q)$ for all voters $i \in \mathcal{N}_{\boldsymbol{B}}$. This implies that (and for resolute rules is equivalent to) $\mathcal{F}(\boldsymbol{B}, \Gamma) = \{J[p \leftrightarrow q] \mid J \in \mathcal{F}(\boldsymbol{B}, \Gamma)\}$, whenever $B_i(p) = B_i(q)$ for all voters $i \in \mathcal{N}_{\boldsymbol{B}}$. Using this formulation, we introduce contextuality to the axiom of neutrality as follows.

Axiom 2.2.13 (Contextual Neutrality)**.** We call a liquid aggregation rule $\mathcal{F}$ *contextually neutral*, if for any pair $(\boldsymbol{B}, \Gamma) \in \mathcal{D}_{\mathcal{F}}$ and any propositional letters $p, q \in \Phi^+$, we have that $\mathcal{F}(\boldsymbol{B}, \Gamma[p \leftrightarrow q]) = \{J[p \leftrightarrow q] \mid J \in \mathcal{F}(\boldsymbol{B}, \Gamma)\}$, whenever $B_i(p) = B_i(q)$ for all voters $i \in \mathcal{N}_{\boldsymbol{B}}$ and $(\boldsymbol{B}, \Gamma[p \leftrightarrow q]) \in \mathcal{D}_{\mathcal{F}}$.

Note that in the resolute case, contextual neutrality reduces to the requirement that if two propositional letters p and q are treated identically by all voters, then $\mathcal{F}(\boldsymbol{B}, \Gamma)$ should accept p if and only if $\mathcal{F}(\boldsymbol{B}, \Gamma[p \leftrightarrow q])$ accepts q, and $\mathcal{F}(\boldsymbol{B}, \Gamma)$ should accept $\neg p$ if and only if $\mathcal{F}(\boldsymbol{B}, \Gamma[p \leftrightarrow q])$ accepts $\neg q$.

In the literature, multiple authors have (informally) proposed alterations of the axiom of neutrality or argued in favor of non-neutral rules (e.g., Costantini et al. [2016], List and Pettit [2002], Mongin [2012], Slavkovik [2014]), and Terzopoulou and Endriss [2020] were the first to formally develop a weakening of the axiom of neutrality. Nevertheless, to the best of our knowledge, no explicitly contextual version of the axiom has been proposed. Possibly, this is due to the fact that most of judgment aggregation theory was developed using a slightly different formalism, where complex formulas are considered part of the agenda, and thus an external integrity constraint is not necessary. Since this formalism lacks an explicit object (the integrity constraint) modeling the 'logical context', the notion of contextual neutrality does not come as naturally there.

Example 2.2.14. To illustrate the difference between neutrality and contextual neutrality, we consider some examples of rules that hold none, one, or both of the properties.

The most illustrative example is due to Costantini et al. [2016]. It regards Kemeny's rule (generalized to judgment aggregation), which selects the complete and consistent judgment set that minimizes the sum of individual disagreements.

For $n = 6$ and $m = 5$, let Γ be an integrity constraint that only allows the judgment sets

$$\begin{aligned}
&\{\neg p_1, \neg p_2, \ \ p_3, \neg p_4, \neg p_5\},\\
&\{\ \ p_1, \ \ p_2, \ \ p_3, \neg p_4, \neg p_5\},\\
&\{\neg p_1, \neg p_2, \neg p_3, \ \ p_4, \neg p_5\},\\
&\{\ \ p_1, \ \ p_2, \neg p_3, \ \ p_4, \neg p_5\},\\
&\{\neg p_1, \neg p_2, \neg p_3, \neg p_4, \ \ p_5\},\\
&\{\ \ p_1, \ \ p_2, \neg p_3, \neg p_4, \ \ p_5\},\\
&\{\ \ p_1, \neg p_2, \neg p_3, \neg p_4, \neg p_5\}.
\end{aligned}$$

Let $\boldsymbol{B} \in \mathcal{B}_{\mathcal{N}}^{\Phi}$ be the profile where the first six of these judgment sets are submitted by one voter each. Then Kemeny's rule selects the seventh judgment set (and no other judgment set), since each voter only disagrees with it on two literals, which happens to be the (unique) minimal total number of disagreements possible.

However, while all voters treat p_1 and p_2 identically, p_1 is collectively accepted and p_2 is collectively rejected. Thus Kemeny's rule is not neutral. Costantini et al. [2016] note that they "believe that this indicates a deficiency with this standard formulation of neutrality rather than with [Kemeny's rule], as this standard formulation does not account for the asymmetries in the set of rational ballots induced by the integrity constraint". Indeed, Kemeny's rule does not fundamentally favor one propositional letter over another, but the integrity constraint does.

Kemeny's rule is, on the other hand, contextually neutral. In this particular example, if we swap the names of p_1 and p_2 in the ballots as well as in the integrity constraint, then the rule accepts p_2 and rejects p_1, satisfying contextual neutrality. If we combine contextual neutrality with language independence, we regain some of the power of neutrality without its counter-intuitive effect on asymmetric integrity constraints, as Corollary 2.3.3 will further illustrate.

Examples of rules that are neutral as well as contextually neutral are quota rules (including unanimity) and oligarchies (including dictatorships).

An example of a rule that is neither contextually neutral nor neutral, is any rule that breaks ties between possible collective judgments according to some fixed tie-breaking order. This tie-breaking order is independent of the integrity constraint and favors some propositional letters over others. It is therefore not (contextually) neutral.

Finally, it is difficult to come up with natural rules that are neutral, but not contextually neutral. However, neutrality does not imply contextual neutrality, since a quota-like rule that uses different quotas for different integrity constraints, might accept two identically treated propositional letters p and q for integrity constraint Γ, while rejecting both of them under the same profile for integrity constraint $\Gamma[p \leftrightarrow q]$. This (rather unnatural) rule is neutral, but not contextually neutral. $\triangle$

A related notion to neutrality is the axiom of unbiasedness, which states that an aggregation rule should not favor a propositional letter p over its negation $\neg p$, or the other way around. Again, the most common notion of unbiasedness does not take into account the role of a propositional letter in the integrity constraint, so we introduce a contextual version as well.

For any integrity constraint $\Gamma \in \mathcal{L}_{\text{sat}}$, let $\Gamma[\sim p]$ denote the variant of Γ where all occurrences of propositional letter $p \in \Phi^+$ are replaced by $\neg p$ (which automatically replaces all occurrences of $\neg p$ by $\neg\neg p$, which is equivalent to p). Similarly, for any profile $\boldsymbol{B} \in \mathcal{B}_{\mathcal{N}}^{\Phi}$, let $\boldsymbol{B}[\sim p] = (B_1[\sim p], \ldots, B_n[\sim p])$ denote the profile where

$$B_i[\sim p](q) = \begin{cases} + & \text{if } q = p \text{ and } B_i(p) = -; \\ - & \text{if } q = p \text{ and } B_i(p) = +; \\ B_i(q) & \text{else.} \end{cases}$$

Finally, for any judgment set $J \subseteq \Phi$, let $J[\sim p]$ denote the variant of J where p is replaced by $\neg p$ if $p \in J$, and $\neg p$ is replaced by p if $\neg p \in J$.

Note that if the negation operations are performed on a profile and integrity constraint simultaneously, and the original profile is rational under the integrity constraint (in the sense of Definition 2.2.2), then it remains rational under the new integrity constraint. However, the same is not the case if we only apply the operation to a profile or only to an integrity constraint, since reversing the truth value of a single propositional letter might render the individual judgments in the original profile inconsistent with the new integrity constraint, or the individual judgments in the new profile with the original integrity constraint. In a contextual version of unbiasedness, we should thus apply the negation operation to the profile as well as the integrity constraint.

The axiom of (contextual) unbiasedness for resolute rules requires that the rule accepts a propositional letter p under some profile if and only if the rule rejects p under the profile where all direct voters change their mind on p (i.e., switch from accepting p to rejecting p, or the other way around). However, this formulation does not directly generalize to irresolute rules. Instead, we formulate the axiom as follows, which reduces to the intended interpretation for resolute rules in a similar manner as we described above for the axiom of neutrality.

Axiom 2.2.15 (Unbiasedness)**.** We call a liquid aggregation rule $\mathcal{F}$ *unbiased*, if for any pair $(\boldsymbol{B}, \Gamma) \in \mathcal{D}_{\mathcal{F}}$ and any propositional letter $p \in \Phi^+$, we have that $\mathcal{F}(\boldsymbol{B}[\sim p], \Gamma) = \{J[\sim p] \mid J \in \mathcal{F}(\boldsymbol{B}, \Gamma)\}$, whenever $(\boldsymbol{B}[\sim p], \Gamma) \in \mathcal{D}_{\mathcal{F}}$.

The contextual version of this axiom then reads as follows.

Axiom 2.2.16 (Contextual Unbiasedness)**.** We call a liquid aggregation rule $\mathcal{F}$ *contextually unbiased*, if for any pair $(\boldsymbol{B}, \Gamma) \in \mathcal{D}_{\mathcal{F}}$ and any propositional letter $p \in \Phi^+$, we have that $\mathcal{F}(\boldsymbol{B}[\sim p], \Gamma[\sim p]) = \{J[\sim p] \mid J \in \mathcal{F}(\boldsymbol{B}, \Gamma)\}$, whenever $(\boldsymbol{B}[\sim p], \Gamma[\sim p]) \in \mathcal{D}_{\mathcal{F}}$.

The final axiom of this section, the axiom of independence, states that the collective judgment on one issue should only depend on the individual judgments on that issue, and not on the individual judgments on any other issues. That is, if some voters change their mind on some propositional letters other than $p \in \Phi^+$, then this should not change the collective judgment on p. For resolute rules, this is a straightforward axiom, but it does not easily generalize to irresolute rules. We propose the following interpretation for irresolute rules.

The intuition behind independence is that if we want to decide whether we accept a propositional letter p, we can look at the support for propositional letter p only, and definitively decide whether or not it should be collectively accepted. However, for irresolute rules, we need to further distinguish between possible acceptance and necessary acceptance.

Consider an irresolute rule which possibly (but not necessarily) accepts propositional letter $p \in \Phi^+$. In effect, it thereby creates two different initial segments of possible collective judgments: one in which it accepts p, and one in which it rejects p. But according to the idea behind independence, whether the rule adds some other propositional letter $q \in \Phi^+$ to either of the initial segments, should not depend on the content of the initial segments. Thus q is either added to both initial segments (i.e., necessary acceptance of q); or to neither initial segment (i.e., necessary rejection of q); or each initial segment again splits into two possible collective judgments, one of which contains q and one of which does not (i.e., possible acceptance of q), thus creating a total of four possible collective judgments. And so on for all other propositional letters.

Note that this interpretation of independence only regards propositional letters and not their negations. However, in combination with either unbiasedness, or collective completeness and complement-freeness, the axiom of independence also affects negations of propositional letters: they too must be decided on independently from each other. Further note that we do not require that we only accept a propositional letter p when we reject its negation $\neg p$ (or the other way around), since this is the content of another axiom: collective complement-freeness. Nor do we require that the decisions on a propositional letter p and its negation $\neg p$ are independent, since those judgments should depend on each other if we want to satisfy collective complement-freeness.

In conclusion, an independent rule either necessarily accepts a propositional letter p, or necessarily rejects it, or accepts it in exactly half of the judgment sets (after intersection with Φ^+), where both halves consist of the exact same judgment sets modulo the presence of p (and modulo negated propositional letters). Possible acceptance thus creates a sort of symmetric binary tree of possible collective decisions. This interpretation of independence formalizes to the following axiom. Note that for resolute rules, the axiom simply requires that if two different profiles treat a propositional letter p identically, then p must either be accepted under both profiles or rejected under both profiles.

Axiom 2.2.17 (Independence)**.** We call a liquid aggregation rule $\mathcal{F}$ *independent*, if for all pairs $(\boldsymbol{B},\Gamma),(\boldsymbol{B}',\Gamma)\in\mathcal{D}_{\mathcal{F}}$ with $\mathcal{N}_{\boldsymbol{B}}=\mathcal{N}_{\boldsymbol{B}'}$ and any propositional letter $p\in\Phi^+$, we have that

- if $p\in J$ for all $J\in\mathcal{F}(\boldsymbol{B},\Gamma)$ then $p\in J'$ for all $J'\in\mathcal{F}(\boldsymbol{B}',\Gamma)$; and
- if $p\notin J$ for all $J\in\mathcal{F}(\boldsymbol{B},\Gamma)$ then $p\notin J'$ for all $J'\in\mathcal{F}(\boldsymbol{B}',\Gamma)$; and
- if $p\in J$ for some $J\in\mathcal{F}(\boldsymbol{B},\Gamma)$ and $p\notin J$ for some $J\in\mathcal{F}(\boldsymbol{B},\Gamma)$, then
  - for each $J\in\mathcal{F}(\boldsymbol{B},\Gamma)$ with $p\in J$, there is a $J'\in\mathcal{F}(\boldsymbol{B},\Gamma)$ with $p\notin J'$ such that $(J\setminus\{p\})\cap\Phi^+=J'\cap\Phi^+$; and
  - for each $J'\in\mathcal{F}(\boldsymbol{B},\Gamma)$ with $p\notin J'$, there is a $J\in\mathcal{F}(\boldsymbol{B},\Gamma)$ with $p\in J$ such that $(J\setminus\{p\})\cap\Phi^+=J'\cap\Phi^+$,

whenever $B_i(p)=B_i'(p)$ for all voters $i\in\mathcal{N}_{\boldsymbol{B}}$.

Note that the third bullet point in the axiom of independence does not refer to the profile $\boldsymbol{B}'$. However, by symmetry between $\boldsymbol{B}$ and $\boldsymbol{B}'$ in the axiom, the first two bullet points ensure that if one profile possibly accepts propositional letter p without necessarily accepting it, then the other profile must do so as well. The third bullet point only adds the requirement that possible acceptance (of *positive* agenda items, i.e., propositional letters) must be shaped like a symmetric binary tree, as we argued for above.

Example 2.2.18. We provide a number of examples and counter-examples of independence. The following sets of judgment sets are possible outputs of independent liquid aggregation rules.

$\{\emptyset\}$	(all literals necessarily rejected)
$\{\{p,\neg p\}\}$	(p and $\neg p$ necessarily accepted)
$\{\{p,q,\neg r\}\}$	(p, q and $\neg r$ necessarily accepted)
$\{\{p,q,\neg r\},\{p,\neg r\}\}$	(q possibly accepted)
$\{\{p,q,\neg r\},\{p,\neg q,\neg r\}\}$	(q possibly accepted)
$\{\{p,q,\neg q,\neg r\},\{p,\neg q,\neg r\}\}$	(q possibly accepted)
$\{\{p,q,\neg r\},\{p,\neg r\},\{p,\neg p,\neg r\}\}$	(q possibly accepted)
$\{\{p,q,\neg q\},\{p,\neg p,\neg r\}\}$	(q poss. acc., negations not independent)
$\{\{p,q\},\{p,\neg q\},\{\neg p,q\},\{\neg p,\neg q\}\}$	(p and q possibly accepted)
$\{\{p,q\},\{p\},\{q\},\emptyset\}$	(p and q possibly accepted)

The following sets of judgment sets are *not* possible outputs of independent liquid aggregation rules.

$\{\{p,q\},\{p\},\{q\}\}$	(no set without p and q)
$\{\{p,q,\neg r\},\{\neg p,q,\neg r\},\{\neg p,\neg q,\neg r\}\}$	(no set with p and without q)
$\{\{p,q,\neg r\},\{\neg p,\neg q,r\}\}$	(p, q and r all depend on each other)

Finally, independent liquid aggregation rules which are also collectively complete and complement-free, can only return sets of judgment sets of the following shape. Let J^+ be the set of necessarily accepted propositional letters, J^- the set of necessarily rejected propositional letters, and ζ the set of propositional letters which are possibly accepted and possibly rejected. Then the output of a collectively complete and complement-free, independent rule is

$$\{J^+ \cup \neg J^- \cup C_\zeta(J') \mid J' \in \mathcal{P}(\zeta)\},$$

where $\neg P$ denotes the set of negations of propositional letters in $P \subseteq \Phi^+$, and $C_P(Q) := Q \cup \neg(P \setminus Q)$ denotes the completion of $Q \subseteq \Phi^+$ with respect to the propositional letters in $P \subseteq \Phi^+$. Examples of such sets are the following.

$$\begin{aligned}
&\{\{p, q, \neg r\}\} && (J^+ = \{p, q\}, J^- = \{r\}, \zeta = \emptyset)\\
&\{\{p, q, \neg r\}, \{p, \neg q, \neg r\}\} && (J^+ = \{p\}, J^- = \{r\}, \zeta = \{q\})\\
&\{\{p, q, r\}, \{p, q, \neg r\}, \{p, \neg q, r\}, \{p, \neg q, \neg r\}\} && (J^+ = \{p\}, J^- = \emptyset, \zeta = \{q, r\})
\end{aligned}$$

$\triangle$

2.2.3 Additional Normative Axioms

The axioms in this section are translations of the axioms by which Lamboray [2009b] characterizes the ranked pairs rule (see Section 3.1) in ordinal preference aggregation. We define the axioms and justify their normative content in this section, and translate the characterization result to our setting in Section 3.3.

Three of the following axioms concern the behavior of aggregation rules with respect to the set of literals that are supported by a majority of voters. The strength of support for a literal is measured by the difference between the number of supporters of a literal, and the number of supporters of the literal's negation, i.e., by the *majority margin*. Formally, the majority margin is defined as follows.

Definition 2.2.19 (Majority Margin)**.** For any profile $\boldsymbol{B} \in \mathcal{B}^\Phi$, the *majority margin* of literal $\ell \in \Phi$ is

$$\begin{aligned}
n_{\boldsymbol{B}}(\ell) = {} & |\{i \in \mathcal{N} \mid \mathrm{rep}_{\boldsymbol{B},\ell}(i) \text{ is defined and } B_{\mathrm{rep}_{\boldsymbol{B},\ell}(i)}(\ell) = +\}|\\
& - |\{i \in \mathcal{N} \mid \mathrm{rep}_{\boldsymbol{B},\ell}(i) \text{ is defined and } B_{\mathrm{rep}_{\boldsymbol{B},\ell}(i)}(\ell) = -\}|.
\end{aligned}$$

Note that our notion of majority margins is closely related to the notion of a proxy profile in the sense that each voter whose representative accepts a literal, is considered to support the literal. However, whereas the proxy profile is undefined for cyclic profiles, the majority margin is always defined, but ignores voters who delegate to a cycle. Further note that for any profile $\boldsymbol{B} \in \mathcal{B}^\Phi$ and literal $\ell \in \Phi$, we have $n_{\boldsymbol{B}}(\ell) = -n_{\boldsymbol{B}}(\sim\ell)$, since voters cannot delegate ℓ and $\sim\ell$ to different peers, and a direct vote in favor of ℓ implies a direct vote against $\sim\ell$.

The literals which are accepted by a majority of voters constitute the *majority set* of the profile. We define such a majority set for each possible threshold of what constitutes a 'large enough' majority, as follows. Note that since the majority margin ignores voters who delegate to a cycle, the majority set of a cyclic profile might contain literals with a very low absolute number of supporters, if many voters delegate to a cycle.

Definition 2.2.20 (Qualified Majority Set)**.** Consider any profile $\boldsymbol{B} \in \mathcal{B}^\Phi$ and let $\gamma \in \mathbb{N}_0$. The set

$$M_\gamma = \{\ell \in \Phi \mid n_{\boldsymbol{B}}(\ell) > \gamma\}$$

of all literals that have a majority margin in $\boldsymbol{B}$ that is strictly greater than γ, is called the *γ-qualified majority set* of $\boldsymbol{B}$. In particular, for $\gamma = 0$, we call the set $M = M_0$ the *strict majority set* of $\boldsymbol{B}$.

Our first two axioms state that reinforcing a majority opinion present in a profile should not create any new possible collective decisions. That is, if a profile has some (qualified or strict) majority set, and we introduce additional voters who display the same majority set and possibly generate a majority for additional literals which previously enjoyed neither majority support nor majority rejection, then the set of possible collective decisions should not increase. Thus, the new voters can only remove some possible collective decisions by breaking the collective 'indifference' on some literals. As Lamboray [2009b] puts it, "[t]he property suggests that 'confirming the majority' of a profile should not lead to creating new solutions, which is in line with the idea that reinforcing the majority does not fundamentally change the aggregation problem."

To formalize this idea, we need two auxiliary definitions. Firstly, profile addition is defined as the concatenation of profiles.

Definition 2.2.21 (Profile Addition)**.** For any two liquid aggregation profiles $\boldsymbol{B} = (B_1, \ldots, B_n) \in \mathcal{B}^\Phi$ and $\boldsymbol{B}' = (B'_1, \ldots, B'_{n'}) \in \mathcal{B}^\Phi$, we write $\boldsymbol{B} + \boldsymbol{B}'$ for the concatenated profile $(B_1, \ldots, B_n, B'_1, \ldots, B'_{n'}) \in \mathcal{B}^\Phi$.

Secondly, we introduce profiles which support a specific literal in a minimal way.

Definition 2.2.22 (Minimal ℓ-Profile)**.** For any literal $\ell \in \Phi$ and the set of voters $\mathcal{N} = \{1, 2\}$, the *minimal ℓ-profile* is the profile $\boldsymbol{B}_\ell = (B_1, B_2) \in \mathcal{B}^\Phi_\mathcal{N}$, where $B_1(p) = +$ and $B_2(p) = -$ for all $p \in \Phi^+ \setminus \{\ell, {\sim}\ell\}$, and $B_1(\ell) = B_2(\ell) = +$.

In other words, the minimal ℓ-profile $\boldsymbol{B}_\ell$ is a profile of two voters who disagree on every literal except ℓ (and ${\sim}\ell$), which they both accept (and reject). Thus, the majority margins of the minimal ℓ-profile are always as follows.

Fact 2.2.23. *If $\boldsymbol{B}_\ell$ is the minimal ℓ-profile for $\ell \in \Phi$, then for $\ell' \in \Phi$, we have*

$$n_{\boldsymbol{B}_\ell}(\ell') = \begin{cases} 2 & \textit{if } \ell' = \ell; \\ -2 & \textit{if } \ell' = {\sim}\ell; \\ 0 & \textit{else.} \end{cases}$$

Note that minimal ℓ-profiles are not generally individually consistent. Therefore, our characterization result in Section 3.3 using minimal ℓ-profiles only applies to the universal domain. The characterization of aggregation rules on the rational domain would have to make use of an individually rational notion of minimal support, which would be a non-trivial adaptation of the definition.

We can now define the majority profile of a profile $\boldsymbol{B} \in \mathcal{B}^\Phi$ as follows. If M is the (strict) majority set of $\boldsymbol{B}$, a majority profile $\boldsymbol{B}(M)$ of $\boldsymbol{B}$ is a 'minimal' profile, expressing majority support for the literals in M as well as for (possibly) some additional literals which received neither majority support nor majority rejection in $\boldsymbol{B}$. That is, $\boldsymbol{B}(M)$ expresses a majority for the same literals as $\boldsymbol{B}$ (and possibly breaks some ties of $\boldsymbol{B}$), but expresses this majority by the smallest possible margin assuming acyclicity (i.e., by a majority margin of 2). Formally, we call such majority profiles *strict majority profiles*. Analogously, we can define γ-*qualified majority profiles* for $\gamma \in \mathbb{N}$ by only duplicating the majorities in $\boldsymbol{B}$ which have a majority margin strictly above γ, and (possibly) lending majority support to some of the literals which enjoy a majority margin in $\boldsymbol{B}$ of exactly γ.

In the formal definition below, M_γ is the γ-qualified majority set of profile $\boldsymbol{B}$, thus containing all literals which enjoy a majority margin strictly greater than γ. And ζ contains those literals which enjoy a majority margin of exactly γ and which we wish to promote to majority-supported literals. Note that since we may choose different sets ζ, the majority profile $\boldsymbol{B}(M_\gamma)$ of a profile $\boldsymbol{B}$ is not unique.

Definition 2.2.24 (Majority Profile)**.** For a profile $\boldsymbol{B} \in \mathcal{B}^\Phi$ with γ-qualified (or strict) majority set M_γ, we say that $\boldsymbol{B}(M_\gamma) \in \mathcal{B}^\Phi$ is a γ-*qualified* (or *strict*) *majority profile* of $\boldsymbol{B}$, if $\boldsymbol{B}(M_\gamma) = \sum_{\ell \in M_\gamma \cup \zeta} \boldsymbol{B}_\ell$ for some complement-free set $\zeta \subseteq \{\ell \in \Phi \mid n_{\boldsymbol{B}}(\ell) = \gamma\}$.

Note that we achieve 'minimality' of the majority profile $\boldsymbol{B}(M_\gamma)$ by defining it as a sum of minimal ℓ-profiles, where ℓ ranges over those literals which should receive majority support. The majority margins of $\boldsymbol{B}$ and $\boldsymbol{B}(M_\gamma)$ are related as follows.

Fact 2.2.25. *If $\boldsymbol{B}(M_\gamma)$ is a γ-qualified (or strict) majority profile of profile $\boldsymbol{B} \in \mathcal{B}^\Phi$ over some set $\zeta \subseteq \{\ell \in \Phi \mid n_{\boldsymbol{B}}(\ell) = \gamma\}$, then for $\ell \in \Phi$, we have*

$$n_{\boldsymbol{B}(M_\gamma)}(\ell) = \begin{cases} 2 & \text{if } n_{\boldsymbol{B}}(\ell) > \gamma \text{ or } \ell \in \zeta; \\ -2 & \text{if } n_{\boldsymbol{B}}(\ell) < \gamma \text{ or } {\sim}\ell \in \zeta; \\ 0 & \text{else.} \end{cases}$$

We can finally define our normative axioms formalizing the idea above, where the first axiom only considers strict majority profiles and the second axiom considers γ-qualified majority profiles for any $\gamma \in \mathbb{N}$. To reiterate, the axioms require that if some profile $\boldsymbol{B} \in \mathcal{B}^\Phi$ generates collective decision $\mathcal{F}(\boldsymbol{B}, \Gamma)$, then strengthening the majority opinion M_γ present in $\boldsymbol{B}$ by concatenating the profile with a majority profile $\boldsymbol{B}(M_\gamma)$ (which expresses the same majority opinion M_γ as $\boldsymbol{B}$ and possibly breaks some ties in $\boldsymbol{B}$) leads to a new collective decision $\mathcal{F}(\boldsymbol{B}+\boldsymbol{B}(M_\gamma),\Gamma)$ which is a subset of the original collective decision $\mathcal{F}(\boldsymbol{B},\Gamma)$. In other words, a judgment set $J \subseteq \Phi$ can only be a possible collective judgment under the strengthened profile $\boldsymbol{B}+\boldsymbol{B}(M_\gamma)$ if it already was a possible collective judgment under the original profile $\boldsymbol{B}$.[3]

Axiom 2.2.26 (Weak Majority Profile Consistency[4])**.** We call a liquid aggregation rule $\mathcal{F}$ *weakly majority profile consistent*, if the following holds for any pair $(\boldsymbol{B},\Gamma) \in \mathcal{D}_\mathcal{F}$.

If M is the strict majority set of $\boldsymbol{B}$ and $\boldsymbol{B}(M)$ is any strict majority profile of $\boldsymbol{B}$, then $\mathcal{F}(\boldsymbol{B}+\boldsymbol{B}(M),\Gamma) \subseteq \mathcal{F}(\boldsymbol{B},\Gamma)$, whenever $(\boldsymbol{B}+\boldsymbol{B}(M),\Gamma) \in \mathcal{D}_\mathcal{F}$.

A stronger version of majority profile consistency does not only consider strict majority profiles, but γ-qualified majority profiles for any $\gamma \in \mathbb{N}_0$, and is defined as follows. Note that weak qualified majority profile consistency implies weak majority profile consistency.

Axiom 2.2.27 (Weak Qualified Majority Profile Consistency)**.** We call a liquid aggregation rule $\mathcal{F}$ *weakly qualified majority profile consistent*, if the following holds for any pair $(\boldsymbol{B},\Gamma) \in \mathcal{D}_\mathcal{F}$ and any $\gamma \in \mathbb{N}_0$.

If M_γ is the γ-qualified majority set of $\boldsymbol{B}$ and $\boldsymbol{B}(M_\gamma)$ is any γ-qualified majority profile of $\boldsymbol{B}$, then $\mathcal{F}(\boldsymbol{B}+\boldsymbol{B}(M_\gamma),\Gamma) \subseteq \mathcal{F}(\boldsymbol{B},\Gamma)$, whenever $(\boldsymbol{B}+\boldsymbol{B}(M_\gamma),\Gamma) \in \mathcal{D}_\mathcal{F}$.

Our third axiom is related to Condorcet's famous paradox. In the context of ordinal voting, Condorcet [1785] showed that there are cases in which there is no alternative that defeats every other alternative in a head-to-head contest. If there is such an alternative however, the Condorcet principle states that we must elect this alternative. If a rule always selects the 'Condorcet winner', we say that the rule is Condorcet consistent. Condorcet consistency translates to judgment aggregation by requiring that in case the strict majority set is consistent with the integrity constraint, then this majority opinion should be respected by all possible collective judgments.

Axiom 2.2.28 (Weak Condorcet Consistency)**.** We call a liquid aggregation rule $\mathcal{F}$ *weakly Condorcet consistent*, if for any pair $(\boldsymbol{B},\Gamma) \in \mathcal{D}_\mathcal{F}$, we have $M \subseteq J$ for all $J \in \mathcal{F}(\boldsymbol{B},\Gamma)$, whenever the strict majority set M of $\boldsymbol{B}$ is consistent with Γ.

[3]In line with the definitions of Lamboray [2009b], we call all axioms in this section 'weak'. The strong version of each axiom (except Condorcet consistency) can be obtained by replacing the relevant set-inclusion with an equality.

[4]Lamboray [2009b] uses an even weaker majority profile consistency axiom, which only requires $\mathcal{F}(\boldsymbol{B}+\boldsymbol{B}(M),\Gamma) \subseteq \mathcal{F}(\boldsymbol{B},\Gamma)$ in case the strict majority profile of $\boldsymbol{B}+\boldsymbol{B}(M)$ is inconsistent with Γ. We omit this extra condition in order to align the formulation of weak majority profile consistency with weak qualified majority profile consistency. The two versions of weak majority profile consistency are interchangeable in all theorems and proofs in this thesis.

For the forth axiom, we need the following auxiliary definition.

Definition 2.2.29 (E-Profile)**.** Any profile $\boldsymbol{B}_{\mathrm{E}} \in \mathcal{B}^{\Phi}$ such that $n_{\boldsymbol{B}_{\mathrm{E}}}(\ell) = 0$ for all literals $\ell \in \Phi$, is called an *E-profile*.

An E-profile (or 'equilibrated profile') is a profile which is collectively undecided on every issue. The idea of (the strong version of) the following axiom is that adding a set of collectively undecided voters to an existing profile does not fundamentally change the aggregation problem. The axiom is related to the idea that aggregation rules should only depend on the majority margins of literals. In fact, Debord [1987] shows that an aggregation rule in ordinal preference aggregation only depends on pairwise majority margins if and only if it is (strongly) E-invariant and anonymous. Our version of the axiom reads as follows.

Axiom 2.2.30 (Weak E-Invariance)**.** We call a liquid aggregation rule $\mathcal{F}$ *weakly E-invariant*, if for any pair $(\boldsymbol{B}, \Gamma) \in \mathcal{D}_{\mathcal{F}}$ and any E-profile $\boldsymbol{B}_{\mathrm{E}}$, we have $\mathcal{F}(\boldsymbol{B}, \Gamma) \subseteq \mathcal{F}(\boldsymbol{B} + \boldsymbol{B}_{\mathrm{E}}, \Gamma)$, whenever $(\boldsymbol{B} + \boldsymbol{B}_{\mathrm{E}}, \Gamma) \in \mathcal{D}_{\mathcal{F}}$.

The fifth axiom is defined to relate the behavior of aggregation rules on 'symmetric' profiles (i.e., profiles which are the sum of two identical profiles) to the behavior of their symmetric components. The normative idea behind (the strong version strong of) the axiom is that cloning each voter should not change the outcome of an aggregation rule. The axiom is defined as follows.

Axiom 2.2.31 (Weak Homogeneity[5])**.** We call a liquid aggregation rule $\mathcal{F}$ *weakly homogeneous*, if for any pair $(\boldsymbol{B}, \Gamma) \in \mathcal{D}_{\mathcal{F}}$, we have $\mathcal{F}(\boldsymbol{B}, \Gamma) \subseteq \mathcal{F}(\boldsymbol{B} + \boldsymbol{B}, \Gamma)$, whenever $(\boldsymbol{B} + \boldsymbol{B}, \Gamma) \in \mathcal{D}_{\mathcal{F}}$.

Our last axiom concerns monotonicity: if the support for some literal increases, this should not harm the literal in the outcome of an aggregation rule. Formally, it states that if a possible collective decision J under profile $\boldsymbol{B}$ contains some literal ℓ, then J should still be a possible collective decision when we add a pair of voters to $\boldsymbol{B}$ who only agree on accepting ℓ.

Axiom 2.2.32 (Monotonic Consistency)**.** We call a liquid aggregation rule $\mathcal{F}$ *monotonically consistent*, if for any pair $(\boldsymbol{B}, \Gamma) \in \mathcal{D}_{\mathcal{F}}$, we have

$$\{J \in \mathcal{F}(\boldsymbol{B}, \Gamma) \mid \ell \in J\} \subseteq \mathcal{F}(\boldsymbol{B} + \boldsymbol{B}_{\ell}, \Gamma),$$

whenever $(\boldsymbol{B} + \boldsymbol{B}_{\ell}, \Gamma) \in \mathcal{D}_{\mathcal{F}}$.

This concludes the collection of axioms necessary to characterize the ranked agenda rule in liquid judgment aggregation, which we show in Section 3.3.

[5] Lamboray [2009b] uses an even weaker homogeneity axiom, which only relates $\boldsymbol{B}$ to $\boldsymbol{B}+\boldsymbol{B}$ for profiles $\boldsymbol{B}$ with an odd number of voters. His axiom is sufficient for the direct democracy setting, but we need this stronger version in our characterization result in Section 3.3 to deal with cyclic delegation. Also note that the term '(weak) homogeneity' is sometimes used in the literature to express that an aggregation rule is invariant to multiplying a profile by any natural number (i.e., $\mathcal{F}(\boldsymbol{B}, \Gamma) \subseteq \mathcal{F}(k \cdot \boldsymbol{B}, \Gamma)$ for any $k \in \mathbb{N}$) instead of just doubling a profile (i.e., $\mathcal{F}(\boldsymbol{B}, \Gamma) \subseteq \mathcal{F}(\boldsymbol{B} + \boldsymbol{B}, \Gamma)$). For our purposes, we only need the latter, weaker notion.

2.3 Three Impossibility Results

In Section 2.1, we described how the proxy profile can be used to translate liquid aggregation problems to classical judgment aggregation problems. However, it is important to note that rational liquid aggregation ballots (in the sense of Definition 2.2.2) need not give rise to rational classical judgment aggregation ballots, since the judgments of a voter's delegates might be inconsistent with the voter's own judgments. It is generally assumed in classical judgment aggregation that judgments of rational voters are individually consistent. Therefore, possibility results in classical judgment aggregation that assume individual rationality, generally do not extend to the rational domain of liquid judgment aggregation when we make use of proxy profiles.

In fact, Grandi and Endriss [2013] show that in the setting of direct democracy, the only static rules that guarantee a consistent outcome under every integrity constraint, are rules that copy the ballot of at least one voter for every given profile (but this voter does not need to be the same voter for each profile). Since proxy profiles of rational liquid profiles need not remain rational, we have the following proposition, forbidding static, complete and consistent rules for the acyclic rational domain.

Proposition 2.3.1. *If $n > 1$ and $m > 1$, then no liquid aggregation rule on the acyclic rational domain is static and collectively complete and consistent.*

Proof. Suppose there exists such a rule $\mathcal{F}$. Fix some $i, j \in \mathcal{N}$ and $p, q \in \Phi^+$ such that $i \neq j$ and $p \neq q$. Consider a profile $\boldsymbol{B}$ where

$$\begin{aligned} B_i(p) &= +, \quad & B_i(q) &= j, \\ B_j(p) &= i, \quad & B_j(q) &= +, \\ B_k(p) &= i, \quad & B_k(q) &= i \end{aligned}$$

for all $k \in \mathcal{N} \setminus \{i, j\}$. This profile is acyclic and rational for the integrity constraint $\Gamma_1 = p \wedge q$. Therefore, by collective completeness and consistency on the acyclic rational domain, we have $\{p, q\} \subseteq J$ for any $J \in \mathcal{F}(\boldsymbol{B}, \Gamma_1)$. Simultaneously, the profile $\boldsymbol{B}$ is rational for the integrity constraint $\Gamma_2 = \neg(p \wedge q)$. Therefore, by collective consistency on the acyclic rational domain, we have $\{p, q\} \not\subseteq J$ for any $J \in \mathcal{F}(\boldsymbol{B}, \Gamma_2)$. But by staticity, we also have $\mathcal{F}(\boldsymbol{B}, \Gamma_1) = \mathcal{F}(\boldsymbol{B}, \Gamma_2)$, which contradicts the above, since $\mathcal{F}(\boldsymbol{B}, \Gamma_1)$ and $\mathcal{F}(\boldsymbol{B}, \Gamma_2)$ are non-empty by definition of a liquid aggregation rule. $\square$

Proposition 2.3.1 implies that if we want to design a liquid aggregation rule that is collectively complete and consistent, and can (at least) be used in case all voters are rational and do not delegate cyclically, then this aggregation rule should explicitly depend on the integrity constraint. In other words, a reasonable liquid aggregation mechanism that can be used for all possible (satisfiable) integrity constraints, cannot behave identically for all those integrity constraints. Examples of collectively complete and consistent rules in classical judgment aggregation which are not static, are Kemeny's rule and the ranked pairs rule (see Example 2.2.14 and Section 3.1).

Another famous result (due to List and Pettit [2002]) states that even if we drop staticity, no reasonable rule guarantees a complete and consistent outcome (where 'reasonable' is now understood to mean anonymous, neutral, unbiased and independent). Since liquid aggregation is a generalization of classical binary aggregation, this result generalizes to the following proposition. Its proof was inspired by the presentation of the proof by Endriss [2016].

Proposition 2.3.2. *If $n > 1$ and $m > 2$, then no liquid aggregation rule on the direct rational domain is anonymous, neutral, unbiased, independent, and collectively complete and consistent.*

Proof. Assume that there is such a rule $\mathcal{F}$. We will show that this leads to a contradiction.

Suppose $n = |\mathcal{N}|$ is odd. Let $p, q, r \in \Phi^+$ be three different propositional letters and consider the integrity constraint $\Gamma = ((p \wedge q) \rightarrow \neg r) \wedge (\neg r \rightarrow (p \vee q))$. Let $\boldsymbol{B} \in \mathcal{B}^{\Phi}_{\mathcal{N}}$ be a rational profile in which all voters submit direct votes only, and

- $\frac{n-1}{2}$ voters accept p, accept q and reject r;
- one voter accepts p, rejects q and accepts r;
- one voter rejects p, accepts q and accepts r;
- $\frac{n-3}{2}$ voters reject p, reject q and accept r.

Then exactly $\frac{n+1}{2}$ voters accept p, exactly $\frac{n+1}{2}$ voters accept q, and exactly $\frac{n+1}{2}$ voters accept r.

By anonymity and independence, the social decision (i.e., necessary acceptance, necessary rejection, or neither) on any individual propositional letter may only depend on the number of voters accepting it (in case of a profile of direct votes). Furthermore, by neutrality, if this number is the same for two different propositional letters, the social decision on both propositional letters must be the same.

We conclude that p, q and r must either all be necessarily accepted, or all be necessarily rejected (in which case $\neg p$, $\neg q$ and $\neg r$ must be necessarily accepted by collective completeness), or all be neither necessarily accepted nor necessarily rejected (in which case $\neg p$, $\neg q$ and $\neg r$ must be neither necessarily rejected nor necessarily accepted by collective completeness and consistency). The first two cases contradict collective consistency, since these directly imply that any possible collective decision is inconsistent with Γ.

But if p, q and r are all neither necessarily accepted nor necessarily rejected, then independence (together with collective completeness and consistency) requires that any complete and complement-free combination of p, q, r and their negations are (subsets of) possible collective judgments (see Example 2.2.18). In particular, $\{p, q, r\}$ must be (a subset of) a possible collective judgment, which contradicts collective consistency again.

Suppose $n = |\mathcal{N}|$ is even. Let $p, q \in \Phi^+$ be two different propositional letters and consider integrity constraint $\Gamma = p \leftrightarrow \neg q$. Let $\boldsymbol{B} \in \mathcal{B}^{\Phi}_{\mathcal{N}}$ be an individually rational profile in which all voters submit direct votes only, and $\frac{n}{2}$ voters accept p and reject q, while the other $\frac{n}{2}$ voters reject p and accept q. By a similar argument using anonymity, independence and unbiasedness (instead of neutrality), we must collectively accept both p and q (or neither) in at least one possible collective judgment, contradicting collective consistency (or completeness). $\square$

Proposition 2.3.2 is a standard result in the field of judgment aggregation for resolute aggregation rules. However, there is (to the best of our knowledge) no standard generalization of the axioms listed in the proposition to irresolute aggregation rules. In particular, the axiom of independence does not generalize to irresolute rules in an obvious way, as we discussed in Section 2.2.2. Therefore, it should be noted that the validity of the result relies strongly on the exact formulation of its axioms, and especially so for the axiom of independence.

Furthermore, the proof of Proposition 2.3.2 relies on the axiom of neutrality and unbiasedness in a strong manner: if a profile treats two literals identically, then the aggregation rule should treat them identically, regardless of their role in the integrity constraint. If we instead consider contextual neutrality and contextual unbiasedness, the proof does not directly go through. However, the following corollary shows that language-independence can recover the impossibility result.

Corollary 2.3.3. *If $n > 1$ and $m > 2$, then no liquid aggregation rule on the direct rational domain is language-independent, anonymous, contextually neutral, contextually unbiased, independent, and collectively complete and consistent.*

Proof. Consider the same case as in the proof of Proposition 2.3.2 for odd n. For this particular (very 'symmetric') integrity constraint Γ, the conditions in the axioms of neutrality and contextual neutrality become equivalent if we assume language independence: since Γ is logically equivalent to $\Gamma[p \leftrightarrow q]$, $\Gamma[p \leftrightarrow r]$ and $\Gamma[q \leftrightarrow r]$,[6] the judgment set $\mathcal{F}(\boldsymbol{B}, \Gamma)$ must be equal to $\mathcal{F}(\boldsymbol{B}, \Gamma[p \leftrightarrow q])$, $\mathcal{F}(\boldsymbol{B}, \Gamma[p \leftrightarrow r])$ and $\mathcal{F}(\boldsymbol{B}, \Gamma[q \leftrightarrow r])$ for any aggregation rule $\mathcal{F}$ by language independence. Thus for any $x, y \in \{p, q, r\}$, the requirement of contextual neutrality that

$$\mathcal{F}(\boldsymbol{B}, \Gamma[x \leftrightarrow y]) = \{J[x \leftrightarrow y] \mid J \in \mathcal{F}(\boldsymbol{B}, \Gamma)\}$$

is equivalent to the requirement of neutrality (in its alternative formulation, see below Axiom 2.2.12) that

$$\mathcal{F}(\boldsymbol{B}, \Gamma) = \{J[x \leftrightarrow y] \mid J \in \mathcal{F}(\boldsymbol{B}, \Gamma)\}.$$

Therefore, we can use the same argument as in the proof of Proposition 2.3.2 to arrive at a contradiction.

For even n, we can also use the argument in the proof of Proposition 2.3.2, since for this integrity constraint Γ, we have an equivalence between Γ and $\Gamma[p \leftrightarrow q]$. $\square$

[6] The equivalences are easier to see, when we rewrite Γ as $(\neg p \vee \neg q \vee \neg r) \wedge (p \vee q \vee r)$.

We conclude from Corollary 2.3.3 that if we wish to design a liquid aggregation mechanism that guarantees collective completeness and consistency, and can (at least) be used in case voters are rational and do not delegate, then we must drop or relax language independence, anonymity, contextual neutrality, contextual unbiasedness or independence.

Chapter 3

The Ranked Agenda Rule

In Chapter 2, we concluded that we must let go of some desirable properties of aggregation rules if we hope to find a collectively complete and consistent liquid aggregation rule. In this chapter, we study a liquid aggregation rule which achieves collective completeness and consistency by violating staticity and independence: the ranked agenda rule. The ranked agenda rule generalizes the ranked pairs rule of Tideman [1987] to the judgment aggregation setting.

In Section 3.1, we define the ranked agenda rule as an algorithm and study its main normative properties. In Section 3.2, we discuss an alternative, functional definition of the rule. And in Section 3.3, we axiomatically characterize the rule.

3.1 The Original Definition

A normatively appealing method of aggregation is finding the collective decision that minimizes overall disagreement, based on the ballots submitted. Kemeny [1959] proposed such a rule for ordinal preference aggregation. Translated to judgment aggregation, Kemeny's rule (which we briefly discussed in Example 2.2.14) selects the judgment set that maximizes the number of voters that agree with the collective judgment on an issue, summed over all issues. Though normatively appealing, Bartholdi et al. [1989] show that finding such optimal sets in ordinal preference aggregation is NP-hard. Since ordinal preference aggregation can be embedded into judgment aggregation (see, e.g., the work of Endriss [2016]), the same holds for our context.

However, the ranked pairs rule proposed by Tideman [1987] can, in the case of ordinal aggregation, serve as an efficient (i.e., polynomial time) approximation of Kemeny's rule [Brill and Fischer, 2012, Zavist and Tideman, 1989]. But the ranked pairs rule also has normatively appealing properties in its own right, some of which it does not share with Kemeny's rule [Lamboray, 2009b, Parkes and Xia, 2012, Tideman, 1987]. Applied to judgment aggregation, the *ranked agenda rule* first finds a ranking of all agenda items based on the strength of the majority supporting it: items which are accepted (or rejected) by most voters are prioritized over items on which accepting and rejecting voters are roughly tied. Consequently, the rule accepts (or rejects; whichever the majority of voters prefers) the items by that order, unless accepting (or rejecting) renders collective judgment set inconsistent with the integrity constraint; in the latter case, the item is rejected (or accepted).

Note that this method always yields a complete and consistent judgment set. However, in ranking the literals in the agenda, some ties might have to be broken in case the majorities for multiple different literals are of the same size. In practice, ties are generally broken by a fixed tie-breaking order. Tideman [1987] originally proposed considering all possible tie-breaking orders, and returning as output the set of resulting judgment sets, but Zavist and Tideman [1989] show that this rule does not satisfy 'independence of clones', the initial motivation for designing the rule. Furthermore, Brill and Fischer [2012] show that this irresolute version of the ranked pairs rule in ordinal aggregation is NP-complete.

A similar increase in complexity is present in the context of judgment aggregation: the ranked agenda rule with a fixed tie-breaking order is Δ_2^{P}-complete (i.e., P^{NP}-complete), while the irresolute version is Σ_2^{P}-complete (i.e., $\mathrm{NP}^{\mathrm{NP}}$-complete) [Endriss and De Haan, 2015]. That is, when given access to a constant-time SAT-solving oracle (which in reality runs in NP-time in the size of the integrity constraint and the number of agenda items), the resolute version can be computed in polynomial time, while the irresolute version cannot be computed in polynomial time (unless P equals NP), but correct solutions can be verified in polynomial time.

A straightforward generalization of the irresolute ranked pairs rule to liquid democracy can be defined as follows.

Definition 3.1.1 (Ranked Agenda Rule)**.** For any profile $\boldsymbol{B} \in \mathcal{B}^{\Phi}$, let $\succeq_{\boldsymbol{B}}$ be the binary relation on Φ where for $\ell_1, \ell_2 \in \Phi$, we have $\ell_1 \succeq_{\boldsymbol{B}} \ell_2$ if and only if $n_{\boldsymbol{B}}(\ell_1) \geq n_{\boldsymbol{B}}(\ell_2)$. Let $\mathcal{LO}(\succeq_{\boldsymbol{B}})$ be the set of linear orders on Φ that are compatible with $\succeq_{\boldsymbol{B}}$ (i.e., the linear orders $\succeq$ where $\ell_1 \succeq \ell_2$ implies $\ell_1 \succeq_{\boldsymbol{B}} \ell_2$). The *ranked agenda rule* is the liquid aggregation rule $\textsc{ra}$ generated by the following process.

- Given profile $\boldsymbol{B}$ and integrity constraint Γ, initialize $\textsc{ra}(\boldsymbol{B}, \Gamma) = \emptyset$.
- For each linear order $\succeq \in \mathcal{LO}(\succeq_{\boldsymbol{B}})$, do the following.
  - Initialize $J = \emptyset$.
  - Iteratively, in the order $\succeq$, consider a literal $\ell \in \Phi$. If $J \cup \{\Gamma\} \models {\sim}\ell$, add ${\sim}\ell$ to J. Otherwise, add ℓ to J.
  - After considering all $\ell \in \Phi$, add J to $\textsc{ra}(\boldsymbol{B}, \Gamma)$.

We can view the different linear orders $\succeq \in \mathcal{LO}(\succeq_{\boldsymbol{B}})$ as refinements of $\succeq_{\boldsymbol{B}}$, where we use different 'tie-breaking orders' τ to break the ties between literals with equal support. Note that we define $\succeq_{\boldsymbol{B}}$ and $\succeq \in \mathcal{LO}(\succeq_{\boldsymbol{B}})$, and therefore the tie-breaking orders τ, on all literals in Φ instead of just on the propositional letters in Φ^+. We do this for a number of reasons. Firstly, it makes the tie-breaking order more expressive: perhaps we would like to treat positive and negative literals equally (thus placing them right after each other in the tie-breaking order), or we would like to accept as many positive literals as possible before accepting negative literals (thus placing all positive literals before all negative literals), or we can define any other complex tie-breaking order. Secondly, the tie-breaking order also resolves ties between a propositional letter and its negation. Thus, the order in which a propositional letter and its negation appear matters, even if they appear directly after each other. Finally, it is a matter of convention: the original ranked pairs rule defines an order on all ordered pairs of candidates, which generalizes to an order on all literals if we embed ordinal aggregation in judgment aggregation in the standard way (see, e.g., the work of Endriss [2016]).

The resolute version of the ranked agenda rule does not consider all possible linear extensions of $\succeq_{\boldsymbol{B}}$, but uses a fixed tie-breaking order τ. It is defined as follows.

Definition 3.1.2 (Ranked Agenda Rule with Tie-Breaking)**.** Let *tie-breaking order* τ be some linear order on Φ. For any profile $\boldsymbol{B} \in \mathcal{B}^\Phi$, let $\succeq_{\boldsymbol{B},\tau}$ be the binary relation on Φ where $\ell_1 \succeq_{\boldsymbol{B},\tau} \ell_2$ if and only if $n_{\boldsymbol{B}}(\ell_1) > n_{\boldsymbol{B}}(\ell_2)$, or $n_{\boldsymbol{B}}(\ell_1) = n_{\boldsymbol{B}}(\ell_2)$ and $(\ell_1, \ell_2) \in \tau$. The *ranked agenda rule with tie-breaker* τ is the liquid aggregation rule $\textsc{ra}_\tau$ generated by the following process.

- Given profile $\boldsymbol{B}$ and integrity constraint Γ, initialize $J = \emptyset$.
- Iteratively, in the order $\succeq_{\boldsymbol{B},\tau}$, consider literal $\ell \in \Phi$. If $J \cup \{\Gamma\} \models {\sim}\ell$, add ${\sim}\ell$ to J. Otherwise, add ℓ to J.
- After considering all $\ell \in \Phi$, return $\textsc{ra}_\tau(\boldsymbol{B}, \Gamma) = \{J\}$.

Note that the ranked agenda rule (with or without tie-breaking) has the universal domain. Further note that since the irresolute ranked agenda rule $\textsc{ra}$ considers every possible tie-breaking order, any resolute version $\textsc{ra}_\tau$ is a *refinement* of $\textsc{ra}$ in the sense that $\textsc{ra}_\tau(\boldsymbol{B}, \Gamma) \subseteq \textsc{ra}(\boldsymbol{B}, \Gamma)$ for any profile $\boldsymbol{B} \in \mathcal{B}^\Phi$ and any integrity constraint $\Gamma \in \mathcal{L}_{\text{sat}}$.

Before we characterize the ranked agenda rule by the axioms from Section 2.2.3, we briefly discuss which classical normative axioms of Section 2.2.2 it satisfies and which it violates. Note that since both $\textsc{ra}$ and $\textsc{ra}_\tau$ have the universal domain, and $\textsc{ra}$ will be characterized on the universal domain, we evaluate its properties as an aggregation rule on the universal domain.

Clearly, $\textsc{ra}$ is irresolute and $\textsc{ra}_\tau$ is resolute, neither are static, and both are language-independent. Furthermore, since both rules consider all literals in the agenda, and only reject them if they render the collective decision inconsistent with the integrity constraint, both rules are collectively complete, complement-free and consistent. And since the majority margin of a literal does not depend on the order of the voters in a profile, both rules are anonymous.

Neither $\textsc{ra}$ nor $\textsc{ra}_\tau$ are neutral (for $m > 1$) by the following counter-example. Let $p, q \in \Phi^+$ be different propositional letters and let $\boldsymbol{B} \in \mathcal{B}^\Phi$

be any profile such that $B_i(p) = B_i(q)$ for all voters $i \in \mathcal{N}_{\boldsymbol{B}}$. Then for integrity constraint $\Gamma = p \wedge \neg q \in \mathcal{L}_{\text{sat}}$, all possible collective decisions under RA and RA_τ contain p and none contain q. This violates neutrality.

The rule RA_τ also violates contextual neutrality, since for the same profile $\boldsymbol{B}$ and integrity constraint $\Gamma = p \leftrightarrow \neg q \in \mathcal{L}_{\text{sat}}$, a tie will have to be broken between accepting p and accepting q. Since Γ is equivalent to $\Gamma[p \leftrightarrow q]$, contextual neutrality (together with resoluteness and language independence) requires that $\text{RA}_\tau(\boldsymbol{B}, \Gamma)$ accepts p if and only if it accepts q, which violates collective consistency.

However, RA is contextually neutral, since whenever the ranked agenda rule under integrity constraint Γ generates possible collective decision J for tie-breaker τ, it generates possible collective decision $J[p \leftrightarrow q]$ under integrity constraint $\Gamma[p \leftrightarrow q]$ for tie-breaker τ', where τ' is a version of τ in which p and q (and $\neg p$ and $\neg q$) switch positions.

Analogously, RA and RA_τ are not unbiased, since under the integrity constraint $\Gamma = p \in \mathcal{L}_{\text{sat}}$, they will never reject p. Moreover, RA_τ is not contextually unbiased, since for a profile $\boldsymbol{B} \in \mathcal{B}^\Phi$ where $n_{\boldsymbol{B}}(p) = n_{\boldsymbol{B}}(\neg p)$, the rule RA_τ will have to break the tie between p and $\neg p$. And RA is contextually unbiased, since for each tie-breaker τ, there is a tie-breaker τ' in which p and $\neg p$ switched positions.

Finally, neither RA nor RA_τ are independent (for $m > 1$) by the following counter-example. Let $p, q \in \Phi^+$ be different propositional letters and assume without loss of generality that τ lists q before p. Consider the integrity constraint $\Gamma = \neg(p \wedge q) \in \mathcal{L}_{\text{sat}}$ and the two profiles $\boldsymbol{B}, \boldsymbol{B}' \in \mathcal{B}^\Phi_{\mathcal{N}}$, where for all voters $i \in \mathcal{N}$,

$$\begin{aligned} B_i(p) &= +, \quad B_i(q) = -, \\ B'_i(p) &= +, \quad B'_i(q) = +. \end{aligned}$$

Then we have $B_i(p) = B'_i(p)$ for all voters $i \in \mathcal{N}$, but RA as well as RA_τ (necessarily) accept p under profile $\boldsymbol{B}$, while they (possibly) reject p under $\boldsymbol{B}'$. This violates independence.

For future reference, we summarize the classical normative properties of the ranked agenda rule in the following proposition. Note that in light of Proposition 2.3.1 and Corollary 2.3.3, the ranked agenda rule violates staticity and independence in order to satisfy collective completeness and consistency, language independence, anonymity, contextual neutrality and contextual unbiasedness.

Proposition 3.1.3. *The ranked agenda rule on the universal domain is language-independent, collectively complete, complement-free and consistent, anonymous, contextually neutral, and contextually unbiased. It is not resolute, static, neutral, unbiased, or independent.*

The ranked agenda rule with tie-breaking on the universal domain is resolute, language-independent, collectively complete, complement-free and consistent, and anonymous. It is not static, neutral, contextually neutral, unbiased, contextually unbiased, or independent.

3.2 Ranked Agenda as a Prudent Rule

The ranked agenda rule can also be defined in an alternative way: it selects exactly those judgment sets, which are maximal with respect to some binary relation (the 'DiscriMin' relation $\succeq^{\boldsymbol{B}}_{\text{disc}}$) that is naturally defined using majority margins. The approach was developed by Lamboray [2009b] in the context of ordinal preference aggregation, and is inspired by an axiom of Arrow and Raynaud [1986], which calls a collective linear order *prudent* if it maximizes the minimal majority margin of all pairwise orderings it accepts (which is a negative number, unless the strict majority set is consistent). In other words, a collective linear order which ranks candidate a above candidate b, even though a majority of voters ranks b above a, can only be prudent if every other collective linear order ranks some candidate a' above some candidate b', against which there is an even larger majority. The ranked pairs rule is an example of a prudent rule.

Although multiple prudent rules have been defined (e.g., by Köhler [1978], Tideman [1987] and Lamboray [2007]), Lamboray [2009a] was the first to axiomatically characterize the *prudent rule* which selects *all* prudent linear orders. Consequently, Lamboray [2009b] showed that the ranked pairs rule is a refinement of the prudent rule,[1] which can be characterized by the same axioms as the prudent rule, supplemented by a monotonicity constraint. In this section, we will generalize the prudent rule and the alternative definition of the ranked pairs rule to liquid judgment aggregation, in preparation of their axiomatic characterizations in Section 3.3.

Analogously to the prudent rule in ordinal preference aggregation, the prudent rule for liquid judgment aggregation selects those complete and consistent judgment sets, which maximize the minimal support for any accepted literal. Formally, it is defined as follows.

Definition 3.2.1 (Prudent Rule)**.** For any profile $\boldsymbol{B} \in \mathcal{B}^\Phi$ and integrity constraint $\Gamma \in \mathcal{L}_{\text{sat}}$, let $\succeq^{\boldsymbol{B}}_{\min}$ be the binary relation on $\mathcal{J}(\Gamma)$ defined by

$$J \succeq^{\boldsymbol{B}}_{\min} J' \text{ if and only if } \min_{\ell \in J} n_{\boldsymbol{B}}(\ell) \geq \min_{\ell \in J'} n_{\boldsymbol{B}}(\ell)$$

for $J, J' \in \mathcal{J}(\Gamma)$. The *prudent rule* is the liquid aggregation rule $\textsc{pr}$ which maps any profile $\boldsymbol{B} \in \mathcal{B}^\Phi$ and integrity constraint $\Gamma \in \mathcal{L}_{\text{sat}}$ to

$$\textsc{pr}(\boldsymbol{B}, \Gamma) = \{J \in \mathcal{J}(\Gamma) \mid \forall J' \in \mathcal{J}(\Gamma)\colon J \succeq^{\boldsymbol{B}}_{\min} J'\}.$$

Note that for a given profile $\boldsymbol{B} \in \mathcal{B}^\Phi$ and integrity constraint $\Gamma \in \mathcal{L}_{\text{sat}}$, each complete and Γ-consistent judgment set $J \in \mathcal{J}(\Gamma)$ is associated with some fixed value $\min_{\ell \in J} n_{\boldsymbol{B}}(\ell) \in \mathbb{N}$. The prudent rule selects exactly those judgment sets $J \in \mathcal{J}(\Gamma)$ for which this value is maximal. Since a finite set of natural

[1] A well-known further refinement of the ranked pairs rule is the leximax rule, which selects the linear orders which lexicographically maximize agreement with the majority judgment [Everaere et al., 2014, Nehring and Pivato, 2019]. Interestingly, [Endriss et al., 2020] show that the leximax rule belongs to the same complexity class as the ranked agenda rule with tie-breaking, even though it is not resolute.

numbers always has a maximal element, the set of $\succeq_{\min}^{\boldsymbol{B}}$-maximal judgment sets is always non-empty, and thus a valid output for a liquid aggregation rule.

The definition of the prudent rule is not algorithmic, but functional, in the sense that it gives us no explicit method to compute its outcome, but selects exactly those judgment sets which have some mathematical property. By adjusting the binary relation on $\mathcal{J}(\Gamma)$ over which we maximize, we can similarly define the ranked agenda rule functionally. We do so as follows.

Definition 3.2.2 (Functional Ranked Agenda Rule)**.** For any profile $\boldsymbol{B} \in \mathcal{B}^{\Phi}$ and integrity constraint $\Gamma \in \mathcal{L}_{\text{sat}}$, let $\succeq_{\text{disc}}^{\boldsymbol{B}}$ be the binary relation on $\mathcal{J}(\Gamma)$ defined by

$$J \succeq_{\text{disc}}^{\boldsymbol{B}} J' \text{ if and only if } \min_{\ell \in J \setminus J'} n_{\boldsymbol{B}}(\ell) \geq \min_{\ell \in J' \setminus J} n_{\boldsymbol{B}}(\ell)$$

for $J, J' \in \mathcal{J}(\Gamma)$, where the minimum over the empty set is defined to equal ∞. The *functional ranked agenda rule* is the liquid aggregation rule $\textsc{ra}^f$ which maps any profile $\boldsymbol{B} \in \mathcal{B}^{\Phi}$ and integrity constraint $\Gamma \in \mathcal{L}_{\text{sat}}$ to

$$\textsc{ra}^f(\boldsymbol{B}, \Gamma) = \{J \in \mathcal{J}(\Gamma) \mid \forall J' \in \mathcal{J}(\Gamma) \colon J \succeq_{\text{disc}}^{\boldsymbol{B}} J'\}.$$

To clarify the difference between the prudent rule and the functional ranked agenda rule, we compare their underlying relations $\succeq_{\min}^{\boldsymbol{B}}$ and $\succeq_{\text{disc}}^{\boldsymbol{B}}$. The relation $\succeq_{\min}^{\boldsymbol{B}}$ of the prudent rule ranks a judgment set $J \in \mathcal{J}(\Gamma)$ above another judgment set $J' \in \mathcal{J}(\Gamma)$ if and only if the minimal support over *all* literals $\ell \in J$ is larger than the minimal support over *all* literals $\ell' \in J'$. On the other hand, the relation $\succeq_{\text{disc}}^{\boldsymbol{B}}$ of the functional ranked agenda rule ranks J above J' if and only if the minimal support over the literals $\ell \in J \setminus J'$ *on which* J *differs with* J' is larger than the minimal support over the literals $\ell' \in J' \setminus J$ *on which* J *differs with* J'. This also explains why the functional ranked agenda rule is a refinement of the ranked agenda rule: if J and J' are tied according to $\succeq_{\min}^{\boldsymbol{B}}$ because of some literal $\ell \in J \cap J'$ which they share, then $\succeq_{\text{disc}}^{\boldsymbol{B}}$ might break that tie by removing ℓ from consideration. But whenever J and J' are tied according to $\succeq_{\text{disc}}^{\boldsymbol{B}}$, they are also tied according to $\succeq_{\min}^{\boldsymbol{B}}$. Thus, $\textsc{ra}^f(\boldsymbol{B}, \Gamma)$ is always a subset of $\textsc{pr}(\boldsymbol{B}, \Gamma)$.

It is not directly clear that the set of $\succeq_{\text{disc}}^{\boldsymbol{B}}$-maximal judgment sets is always non-empty (and thus a valid output for a liquid aggregation rule), since $\succeq_{\text{disc}}^{\boldsymbol{B}}$ need not be transitive (as Lamboray [2009b, pp. 133–134] shows). However, the asymmetric part $\succ_{\text{disc}}^{\boldsymbol{B}}$ of $\succeq_{\text{disc}}^{\boldsymbol{B}}$ is transitive by the following (simple, but tedious) argument.

Proof (transitivity of $\succ_{\text{disc}}^{\boldsymbol{B}}$). Let $J, J', J'' \in \mathcal{J}(\Gamma)$. Choose a literal from $J \setminus J'$ which achieves the minimum $\min_{\ell \in J \setminus J'} n_{\boldsymbol{B}}(\ell)$, and call it $\ell_{J'}^{J}$. Similarly, choose literals $\ell_{J}^{J'}$, $\ell_{J''}^{J}$, etc. Suppose that $J \succ_{\text{disc}}^{\boldsymbol{B}} J'$ and $J' \succ_{\text{disc}}^{\boldsymbol{B}} J''$. In other words,

$$n_{\boldsymbol{B}}(\ell_{J'}^{J}) > n_{\boldsymbol{B}}(\ell_{J}^{J'}) \tag{1}$$

and

$$n_{\boldsymbol{B}}(\ell_{J''}^{J'}) > n_{\boldsymbol{B}}(\ell_{J'}^{J''}). \tag{2}$$

Then to prove transitivity, we must show that we have $J \succ_{\text{disc}}^{\boldsymbol{B}} J''$, i.e., that $n_{\boldsymbol{B}}(\ell_{J''}^{J}) > n_{\boldsymbol{B}}(\ell_{J}^{J''})$. We do so by deriving a contradiction from the opposite assumption,

$$n_{\boldsymbol{B}}(\ell_{J}^{J''}) \geq n_{\boldsymbol{B}}(\ell_{J''}^{J}). \tag{3}$$

Suppose we further have that $\ell_{J''}^{J} \in J'$. Since $\ell_{J''}^{J} \notin J''$, this implies that

$$n_{\boldsymbol{B}}(\ell_{J''}^{J}) \geq \min_{\ell \in J' \setminus J''} n_{\boldsymbol{B}}(\ell) =: n_{\boldsymbol{B}}(\ell_{J''}^{J'}). \tag{4}$$

Therefore, we have

$$\min_{\ell \in J'' \setminus J} n_{\boldsymbol{B}}(\ell) =: n_{\boldsymbol{B}}(\ell_{J}^{J''}) \underset{(3)}{\geq} n_{\boldsymbol{B}}(\ell_{J''}^{J}) \underset{(4)}{\geq} n_{\boldsymbol{B}}(\ell_{J''}^{J'}) \underset{(2)}{>} n_{\boldsymbol{B}}(\ell_{J'}^{J''}).$$

But $\min_{\ell \in J'' \setminus J} n_{\boldsymbol{B}}(\ell) > n_{\boldsymbol{B}}(\ell_{J'}^{J''})$ together with $\ell_{J'}^{J''} \in J''$ implies that $\ell_{J'}^{J''} \in J$. Since $\ell_{J'}^{J''} \notin J'$, this implies

$$n_{\boldsymbol{B}}(\ell_{J'}^{J''}) \geq \min_{\ell \in J \setminus J'} n_{\boldsymbol{B}}(\ell) =: n_{\boldsymbol{B}}(\ell_{J'}^{J}). \tag{5}$$

Therefore, we have

$$\min_{\ell \in J' \setminus J''} n_{\boldsymbol{B}}(\ell) =: n_{\boldsymbol{B}}(\ell_{J''}^{J'}) \underset{(2)}{>} n_{\boldsymbol{B}}(\ell_{J'}^{J''}) \underset{(5)}{\geq} n_{\boldsymbol{B}}(\ell_{J'}^{J}) \underset{(1)}{>} n_{\boldsymbol{B}}(\ell_{J}^{J'}).$$

But $\min_{\ell \in J' \setminus J''} n_{\boldsymbol{B}}(\ell) > n_{\boldsymbol{B}}(\ell_{J}^{J'})$ together with $\ell_{J}^{J'} \in J'$ implies that $\ell_{J}^{J'} \in J''$. Since $\ell_{J}^{J'} \notin J$, this implies

$$n_{\boldsymbol{B}}(\ell_{J}^{J'}) \geq \min_{\ell \in J'' \setminus J} n_{\boldsymbol{B}}(\ell) =: n_{\boldsymbol{B}}(\ell_{J}^{J''}). \tag{6}$$

Therefore, we have

$$\begin{aligned} n_{\boldsymbol{B}}(\ell_{J}^{J'}) &\underset{(6)}{\geq} n_{\boldsymbol{B}}(\ell_{J}^{J''}) \underset{(3)}{\geq} n_{\boldsymbol{B}}(\ell_{J''}^{J}) \underset{(4)}{\geq} n_{\boldsymbol{B}}(\ell_{J''}^{J'}) \\ &\underset{(2)}{>} n_{\boldsymbol{B}}(\ell_{J'}^{J''}) \underset{(5)}{\geq} n_{\boldsymbol{B}}(\ell_{J'}^{J}) \underset{(1)}{>} n_{\boldsymbol{B}}(\ell_{J}^{J'}). \end{aligned}$$

But $n_{\boldsymbol{B}}(\ell_{J}^{J'}) > n_{\boldsymbol{B}}(\ell_{J}^{J'})$ is a contradiction. We conclude that $\ell_{J''}^{J} \notin J'$. Since we also have $\ell_{J''}^{J} \in J$, this means that

$$n_{\boldsymbol{B}}(\ell_{J''}^{J}) \geq \min_{\ell \in J \setminus J'} n_{\boldsymbol{B}}(\ell) =: n_{\boldsymbol{B}}(\ell_{J'}^{J}).$$

By analogous arguments, we can derive the inequalities $n_{\boldsymbol{B}}(\ell_{J}^{J'}) \geq n_{\boldsymbol{B}}(\ell_{J''}^{J'})$ and $n_{\boldsymbol{B}}(\ell_{J'}^{J''}) \geq n_{\boldsymbol{B}}(\ell_{J}^{J''})$. Therefore, we have

$$\begin{aligned} n_{\boldsymbol{B}}(\ell_{J''}^{J}) &\geq n_{\boldsymbol{B}}(\ell_{J'}^{J}) \underset{(1)}{>} n_{\boldsymbol{B}}(\ell_{J}^{J'}) \geq n_{\boldsymbol{B}}(\ell_{J''}^{J'}) \\ &\underset{(2)}{>} n_{\boldsymbol{B}}(\ell_{J'}^{J''}) \geq n_{\boldsymbol{B}}(\ell_{J}^{J''}) \underset{(3)}{\geq} n_{\boldsymbol{B}}(\ell_{J''}^{J}). \end{aligned}$$

Thus, we have arrived at $n_{\boldsymbol{B}}(\ell_{J''}^{J}) > n_{\boldsymbol{B}}(\ell_{J''}^{J})$, our final contradiction. $\square$

We have shown that $\succ^{\boldsymbol{B}}_{\text{disc}}$ is transitive. Since $\succeq^{\boldsymbol{B}}_{\text{disc}}$ is also complete, the set of $\succeq^{\boldsymbol{B}}_{\text{disc}}$-maximal judgment sets is always non-empty, and thus a valid output of the liquid aggregation rule $\textsc{ra}^f$. What remains to be shown, is that the functional definition of the ranked agenda rule is indeed equivalent to its algorithmic definition.

To do so, let $[\![\succeq_{\boldsymbol{B}}]\!]_\sim = \{E_1, \dots, E_k\}$ denote the set of equivalence classes (in decreasing order) under the majority relation $\succeq_{\boldsymbol{B}}$ of $\boldsymbol{B} \in \mathcal{B}^\Phi$ on Φ, as defined in Definition 3.1.1 of the algorithmic ranked agenda rule. That is, equivalence class $E_1 \in [\![\succeq_{\boldsymbol{B}}]\!]_\sim$ contains the literals which receive the largest support in $\boldsymbol{B}$, equivalence class $E_2 \in [\![\succeq_{\boldsymbol{B}}]\!]_\sim$ contains the literals which receive the next largest support, etc. The algorithmic ranked agenda rule can then be seen as an algorithm which nondeterministically breaks ties between all elements of $E_1 \in [\![\succeq_{\boldsymbol{B}}]\!]_\sim$, then breaks ties between all elements of $E_2 \in [\![\succeq_{\boldsymbol{B}}]\!]_\sim$, etc. For any judgment set $J \subseteq \Phi$ and $i \in \{1, \dots, k\}$, let $E_i(J) = (E_1 \cup \dots \cup E_i) \cap J$ denote the set of literals in J which the ranked agenda rule considers before or at stage i. Using this notation, the following lemma states that the algorithmic ranked agenda rule accepts as many elements as possible of any $E_i \in [\![\succeq_{\boldsymbol{B}}]\!]_\sim$, before moving on to $E_{i+1} \in [\![\succeq_{\boldsymbol{B}}]\!]_\sim$.

Lemma 3.2.3. *Let $\boldsymbol{B} \in \mathcal{B}^\Phi$ and $\Gamma \in \mathcal{L}_{sat}$. For any judgment set $J \subseteq \Phi$, we have $J \in \textsc{ra}(\boldsymbol{B}, \Gamma)$ if and only if J is consistent with Γ and $E_i(J) \cup \{\ell\}$ is inconsistent with Γ for all $\ell \in E_i \setminus J$ and all $E_i \in [\![\succeq_{\boldsymbol{B}}]\!]_\sim$.*

Proof. Suppose $J \in \textsc{ra}(\boldsymbol{B}, \Gamma)$, and assume that there are $E_i \in [\![\succeq_{\boldsymbol{B}}]\!]_\sim$ and $\ell \in E_i \setminus J$ such that $E_i(J) \cup \{\ell\}$ is consistent with Γ. Let $\succeq \in \mathcal{LO}(\succeq_{\boldsymbol{B}})$ be the linear order witnessing that $J \in \textsc{ra}(\boldsymbol{B}, \Gamma)$. When the ranked agenda algorithm arrives at the $\succeq$-largest literal of E_i, it has already selected $E_{i-1}(J)$. It then continues to select literals of E_i in $\succeq$-order. When it arrives at ℓ, it has selected a subset of $E_i(J)$ (which by assumption remains consistent with Γ if we add ℓ) and thus it accepts ℓ. But then $\ell \in J$, which contradicts our assumption that $\ell \in E_i \setminus J$. We conclude that $E_i(J) \cup \{\ell\}$ is inconsistent with Γ for all $E_i \in [\![\succeq_{\boldsymbol{B}}]\!]_\sim$ and all $\ell \in E_i \setminus J$. Furthermore, since the ranked agenda algorithm only accepts literals if it does not render the judgment set inconsistent with Γ, judgment set J is consistent with Γ.

Contrariwise, suppose $J \subseteq \Phi$ is consistent with Γ, and $E_i(J) \cup \{\ell\}$ is inconsistent with Γ for all $E_i \in [\![\succeq_{\boldsymbol{B}}]\!]_\sim$ and all $\ell \in E_i \setminus J$. Let $\succeq$ be a linear order on Φ which ranks all literals in $E_i \in [\![\succeq_{\boldsymbol{B}}]\!]_\sim$ before all literals in $E_j \in [\![\succeq_{\boldsymbol{B}}]\!]_\sim$ whenever $i < j$ (in other words, $\succeq \in \mathcal{LO}(\succeq_{\boldsymbol{B}})$), and within each $E_i \in [\![\succeq_{\boldsymbol{B}}]\!]_\sim$, it ranks all literals in $E_i \cap J$ before all literals in $E_i \setminus J$. Then $\succeq$ witnesses that $J \in \textsc{ra}(\boldsymbol{B}, \Gamma)$, since within every $E_i \in [\![\succeq_{\boldsymbol{B}}]\!]_\sim$ (in index-increasing order), the ranked agenda algorithm first accepts all literals in $E_i \cap J$ (which is allowed since the literals accepted so far form a subset of $E_i(J) \subseteq J$, and J is consistent with Γ) and then rejects all literals in $E_i \setminus J$ (since it has already selected $E_i(J)$ and $E_i(J) \cup \{\ell\}$ is inconsistent with Γ for all $\ell \in E_i \setminus J$). $\square$

Finally, the following proposition shows that the algorithmic definition and the functional definition of the ranked agenda rule are equivalent.

Proposition 3.2.4. *For any profile $\boldsymbol{B} \in \mathcal{B}^\Phi$ and integrity constraint $\Gamma \in \mathcal{L}_{sat}$, we have*

$$\text{RA}(\boldsymbol{B}, \Gamma) = \text{RA}^f(\boldsymbol{B}, \Gamma).$$

Proof. Fix some profile $\boldsymbol{B} \in \mathcal{B}^\Phi$ and integrity constraint $\Gamma \in \mathcal{L}_{\text{sat}}$.

Let $J \in \text{RA}(\boldsymbol{B}, \Gamma)$ and suppose $J \notin \text{RA}^f(\boldsymbol{B}, \Gamma)$. Then there must be some $J' \in \mathcal{J}(\Gamma)$ such that $J' \succ^{\boldsymbol{B}}_{\text{disc}} J$, i.e., such that

$$\min_{\ell \in J' \setminus J} n_{\boldsymbol{B}}(\ell) > \min_{\ell \in J \setminus J'} n_{\boldsymbol{B}}(\ell).$$

Since J and J' are maximally Γ-consistent and therefore complete, we have $\ell \in J \setminus J'$ if and only if $\sim\!\ell \in J' \setminus J$ for any $\ell \in \Phi$. And since $n_{\boldsymbol{B}}(\ell) = -n_{\boldsymbol{B}}(\sim\!\ell)$ for any $\ell \in \Phi$, we thus have

$$\max_{\ell \in J' \setminus J} n_{\boldsymbol{B}}(\ell) = \max_{\sim\ell \in J \setminus J'} n_{\boldsymbol{B}}(\ell) = \max_{\ell \in J \setminus J'} -n_{\boldsymbol{B}}(\ell) = -\min_{\ell \in J \setminus J'} n_{\boldsymbol{B}}(\ell).$$

With the analogous equality for $\max_{\ell \in J \setminus J'} n_{\boldsymbol{B}}(\ell)$, we obtain

$$\max_{\ell \in J' \setminus J} n_{\boldsymbol{B}}(\ell) > \max_{\ell \in J \setminus J'} n_{\boldsymbol{B}}(\ell).$$

Let $\ell' \in J' \setminus J$ be a witness of the last inequality, and let $E_i \in [\![\succeq_{\boldsymbol{B}}]\!]_\sim$ be its equivalence class (i.e., E_i is the equivalence class with the smallest index on which J and J' differ). We must then have $E_i(J) \subseteq E_i(J')$. And since $\ell' \in E_i \cap J' \subseteq E_i(J')$, we have $E_i(J) \cup \{\ell'\} \subseteq E_i(J')$. But $E_i(J') \subseteq J'$ is consistent with Γ and thus $E_i(J) \cup \{\ell'\}$ is consistent with Γ. By Lemma 3.2.3, we thus have $J \notin \text{RA}(\boldsymbol{B}, \Gamma)$, which contradicts our assumptions. This proves that $\text{RA}(\boldsymbol{B}, \Gamma) \subseteq \text{RA}^f(\boldsymbol{B}, \Gamma)$.

Contrariwise, let $J \in \text{RA}^f(\boldsymbol{B}, \Gamma)$. Then J is consistent with Γ by definition of the rule RA^f. Suppose that $J \notin \text{RA}(\boldsymbol{B}, \Gamma)$. Then by Lemma 3.2.3, there are $E_i \in [\![\succeq_{\boldsymbol{B}}]\!]_\sim$ and $\ell' \in E_i \setminus J$ such that $E_i(J) \cup \{\ell'\}$ is consistent with Γ. Let $\succeq'$ be a linear order on Φ which ranks all literals in $E_i \in [\![\succeq_{\boldsymbol{B}}]\!]_\sim$ before all literals in $E_j \in [\![\succeq_{\boldsymbol{B}}]\!]_\sim$ whenever $i < j$, and within each $E_i \in [\![\succeq_{\boldsymbol{B}}]\!]_\sim$, ranks all literals in $E_i \cap J$ before all literals in $E_i \setminus J$, placing ℓ' directly after $E_i \cap J$. Then $\succeq' \in \mathcal{LO}(\succeq_{\boldsymbol{B}})$ and thus the ranked agenda algorithm applied to it, generates some Γ-consistent $J' \subseteq \Phi$ such that $E_i(J) \cup \{\ell'\} \subseteq E_i(J')$ (i.e., the algorithm adds exactly those literals to J' which are also in J, until it reaches ℓ', which it adds to J' even though $\ell' \notin J$, and then continues considering literals from equivalence classes E_j for $j \geq i$). But then $\ell' \in J' \setminus J$, and all literals which are in J but are not added to J' (i.e., the literals in $J \setminus J'$) are in equivalence classes E_j with $j > i$. Thus, $\ell' \in J' \setminus J$ has a strictly larger majority margin than any literal in $J \setminus J'$. In other words, we have

$$\max_{\ell \in J' \setminus J} n_{\boldsymbol{B}}(\ell) = n_{\boldsymbol{B}}(\ell') > \max_{\ell \in J \setminus J'} n_{\boldsymbol{B}}(\ell).$$

But since J and J' are maximally Γ-consistent and therefore complete, we have $\ell \in J \setminus J'$ if and only if $\sim\!\ell \in J' \setminus J$ for any $\ell \in \Phi$. Therefore, since

$n_{\boldsymbol{B}}(\ell) = -n_{\boldsymbol{B}}(\sim\ell)$ for any $\ell \in \Phi$, we have

$$\min_{\ell \in J' \setminus J} n_{\boldsymbol{B}}(\ell) = -\max_{\sim\ell \in J \setminus J'} n_{\boldsymbol{B}}(\sim\ell) > -\max_{\sim\ell \in J' \setminus J} n_{\boldsymbol{B}}(\sim\ell) = \min_{\ell \in J \setminus J'} n_{\boldsymbol{B}}(\ell)$$

and thus $J' \succ^{\boldsymbol{B}}_{\text{disc}} J$. We conclude that $J \notin \textsc{ra}^f(\boldsymbol{B}, \Gamma)$, contradicting to our assumptions. This proves that $\textsc{ra}^f(\boldsymbol{B}, \Gamma) \subseteq \textsc{ra}(\boldsymbol{B}, \Gamma)$. □

3.3 A Prudent Characterization

The functional definition of the ranked agenda rule allows us to axiomatically characterize it as a refinement of the prudent rule. In this section, we generalize the characterization results of the prudent rule and the ranked pairs rule by Lamboray [2009a,b] from ordinal preference aggregation in direct democracy to judgment aggregation in liquid democracy. The proofs by Lamboray [2009a,b] rely mostly on the maximality of possible collective decisions under relations $\succeq^{\boldsymbol{B}}_{\text{min}}$ and $\succeq^{\boldsymbol{B}}_{\text{disc}}$, and on the analogues of Lemma 3.2.3 and Proposition 3.2.4, and rarely use properties of aggregation rules or judgment sets that are unique to the ordinal preference aggregation setting. Therefore, our proofs stay very close to the original proofs for the restricted setting.

The following proposition lists the axiomatic properties of the prudent rule, as defined in Section 2.2.3.

Proposition 3.3.1. *The prudent rule on the universal domain is collectively complete and consistent, weakly majority profile consistent, weakly Condorcet consistent, weakly E-invariant, and weakly homogeneous.*

Proof. Fix some profile $\boldsymbol{B} \in \mathcal{B}^\Phi$ and integrity constraint $\Gamma \in \mathcal{L}_{\text{sat}}$.

Collective completeness and consistency are immediate from the definition of the prudent rule.

For weak majority profile consistency, let M be the strict majority set of $\boldsymbol{B}$ and let $\boldsymbol{B}(M)$ be a strict majority profile of $\boldsymbol{B}$. By definition of the prudent rule, there are fixed values $N_{\boldsymbol{B}} \in \mathbb{N}$ and $N_{\boldsymbol{B}+\boldsymbol{B}(M)} \in \mathbb{N}$ such that for $J \in \mathcal{J}(\Gamma)$, we have $J \in \textsc{pr}(\boldsymbol{B}, \Gamma)$ if and only if $\min_{\ell \in J} n_{\boldsymbol{B}+\boldsymbol{B}(M)}(\ell) = N_{\boldsymbol{B}}$, and $J \in \textsc{pr}(\boldsymbol{B} + \boldsymbol{B}(M), \Gamma)$ if and only if $\min_{\ell \in J} n_{\boldsymbol{B}}(\ell) = N_{\boldsymbol{B}+\boldsymbol{B}(M)}$. Moreover, for all complete and Γ-consistent judgment sets $J \in \mathcal{J}(\Gamma)$, we have that $\min_{\ell \in J} n_{\boldsymbol{B}}(\ell) \leq N_{\boldsymbol{B}}$ and $\min_{\ell \in J} n_{\boldsymbol{B}+\boldsymbol{B}(M)}(\ell) \leq N_{\boldsymbol{B}+\boldsymbol{B}(M)}$.

We make a case distinction.

- Suppose $N_{\boldsymbol{B}} > 0$. Then $\textsc{pr}(\boldsymbol{B}, \Gamma)$ only accepts literals $\ell \in \Phi$ which enjoy strict majority support. But this implies that all $\sim\ell$ for $\ell \in M$ are rejected. Since any $J \in \textsc{pr}(\boldsymbol{B}, \Gamma)$ is complete, we must thus have $\textsc{pr}(\boldsymbol{B}, \Gamma) = \{M\}$.

  But then M must be a complete judgment set and thus there are no literals with a majority margin of exactly 0. Thus Fact 2.2.25 implies that the literals which receive majority support in $\boldsymbol{B}$, receive even stronger support in $\boldsymbol{B} + \boldsymbol{B}(M)$, and the literals which do not receive majority support in $\boldsymbol{B}$, receive even weaker support in $\boldsymbol{B} + \boldsymbol{B}(M)$. Therefore, by the same argument as for $\textsc{pr}(\boldsymbol{B}, \Gamma)$, we have $\textsc{pr}(\boldsymbol{B} + \boldsymbol{B}(M), \Gamma) = \{M\}$.

We conclude that $\textsc{pr}(\boldsymbol{B} + \boldsymbol{B}(M), \Gamma) \subseteq \textsc{pr}(\boldsymbol{B}, \Gamma)$.

- Suppose $N_{\boldsymbol{B}} \leq 0$. Let $J^* \in \textsc{pr}(\boldsymbol{B}, \Gamma)$. Then for all literals $\ell \in J^*$, we have $n_{\boldsymbol{B}}(\ell) \geq N_{\boldsymbol{B}}$. By Fact 2.2.25, any literal $\ell \in \Phi$ with $n_{\boldsymbol{B}}(\ell) \geq N_{\boldsymbol{B}}$ has $n_{\boldsymbol{B}+\boldsymbol{B}(M)}(\ell) \geq N_{\boldsymbol{B}} - 2$. Therefore, $\min_{\ell \in J^*} n_{\boldsymbol{B}+\boldsymbol{B}(M)}(\ell) \geq N_{\boldsymbol{B}} - 2$.

  Let $J \in \textsc{pr}(\boldsymbol{B} + \boldsymbol{B}(M), \Gamma)$. Then $J \succeq_{\min}^{\boldsymbol{B}+\boldsymbol{B}(M)} J^*$ by definition of the prudent rule. Thus, we have $\min_{\ell \in J} n_{\boldsymbol{B}+\boldsymbol{B}(M)}(\ell) \geq N_{\boldsymbol{B}} - 2$. But by Fact 2.2.25, any literal $\ell \in \Phi$ with (negative) majority margin $n_{\boldsymbol{B}}(\ell) < N_{\boldsymbol{B}}$ has $n_{\boldsymbol{B}+\boldsymbol{B}(M)}(\ell) < N_{\boldsymbol{B}} - 2$. Therefore, $\min_{\ell \in J} n_{\boldsymbol{B}}(\ell) \geq N_{\boldsymbol{B}}$. Thus, we have $\min_{\ell \in J} n_{\boldsymbol{B}}(\ell) = N_{\boldsymbol{B}}$ and therefore $J \in \textsc{pr}(\boldsymbol{B}, \Gamma)$.

  We conclude that $\textsc{pr}(\boldsymbol{B} + \boldsymbol{B}(M), \Gamma) \subseteq \textsc{pr}(\boldsymbol{B}, \Gamma)$.

For weak Condorcet consistency, let M be the strict majority set of $\boldsymbol{B}$ and suppose M is consistent with Γ. Then for any $\ell \in M$, we have $n_{\boldsymbol{B}}(\ell) > 0$ and thus $n_{\boldsymbol{B}}(\sim\ell) < 0$. And for any $\ell' \in \Phi$ with $\ell' \notin M$ and $\sim\ell' \notin M$, we have $n_{\boldsymbol{B}}(\ell') \leq 0$ and $n_{\boldsymbol{B}}(\sim\ell') \leq 0$, and therefore $n_{\boldsymbol{B}}(\ell') = n_{\boldsymbol{B}}(\sim\ell') = 0$. Thus, since M is consistent with Γ, there exists some complete and Γ-consistent judgment set $J^* \in \mathcal{J}(\Gamma)$ such that $M \subseteq J^*$ and $\min_{\ell \in J^*} n_{\boldsymbol{B}}(\ell) \geq 0$.

But then for any $J \in \textsc{pr}(\boldsymbol{B}, \Gamma)$, we have that $J \succeq_{\min}^{\boldsymbol{B}} J^*$ and therefore $\min_{\ell \in J} n_{\boldsymbol{B}}(\ell) \geq 0$. Since J is complete and all $\ell \in M$ have $n_{\boldsymbol{B}}(\sim\ell) < 0$, we conclude that $M \subseteq J$.

For weak E-invariance, let $\boldsymbol{B}_{\mathrm{E}}$ be an E-profile. Then clearly, we have $\min_{\ell \in J} n_{\boldsymbol{B}}(\ell) = \min_{\ell \in J} n_{\boldsymbol{B}+\boldsymbol{B}_{\mathrm{E}}}(\ell)$ for all $J \in \mathcal{J}(\Gamma)$. Therefore, if $J \in \textsc{pr}(\boldsymbol{B}, \Gamma)$, then $J \in \textsc{pr}(\boldsymbol{B} + \boldsymbol{B}_{\mathrm{E}}, \Gamma)$.

For weak homogeneity, note that $n_{\boldsymbol{B}+\boldsymbol{B}}(\ell) = 2 \cdot n_{\boldsymbol{B}}(\ell)$ for all literals $\ell \in \Phi$. Therefore, $\min_{\ell \in J} n_{\boldsymbol{B}+\boldsymbol{B}}(\ell) = 2 \cdot \min_{\ell \in J} n_{\boldsymbol{B}}(\ell)$ for all $J \in \mathcal{J}(\boldsymbol{B})$. Thus, any $\succeq_{\min}^{\boldsymbol{B}}$-maximal judgment set $J \in \mathcal{J}(\Gamma)$ is also a $\succeq_{\min}^{\boldsymbol{B}+\boldsymbol{B}}$-maximal judgment set. Therefore, $\textsc{pr}(\boldsymbol{B}, \Gamma) \subseteq \textsc{pr}(\boldsymbol{B} + \boldsymbol{B}, \Gamma)$. □

We see that the prudent rule in the liquid aggregation setting satisfies the same properties as the prudent rule in the direct ordinal aggregation setting (see Corollary 16 by Lamboray [2009a]). In fact, the following theorem shows that the prudent rule is the largest liquid aggregation rule that satisfies these properties, in the sense that any other rule $\mathcal{F}$ which satisfies these properties, is a refinement of the prudent rule, i.e., $\mathcal{F}(\boldsymbol{B}, \Gamma) \subseteq \textsc{pr}(\boldsymbol{B}, \Gamma)$ for all profiles $\boldsymbol{B} \in \mathcal{B}^{\Phi}$ and integrity constraints $\Gamma \in \mathcal{L}_{\mathrm{sat}}$.

Theorem 3.3.2. *The prudent rule is the $\subseteq$-largest liquid aggregation rule on the universal domain that is collectively complete and consistent, weakly majority profile consistent, weakly Condorcet consistent, weakly E-invariant, and weakly homogeneous.*

Proof. By Proposition 3.3.1, the prudent rule satisfies all the properties in the theorem. What remains to be shown is that for any liquid aggregation rule $\mathcal{F}$ on the universal domain that satisfies these properties, we have that $\mathcal{F}(\boldsymbol{B},\Gamma) \subseteq \textsc{pr}(\boldsymbol{B},\Gamma)$ for all $\boldsymbol{B} \in \mathcal{B}^\Phi$ and $\Gamma \in \mathcal{L}_{\text{sat}}$. Let $\mathcal{F}$ be such a rule and fix a profile $\boldsymbol{B} \in \mathcal{B}^\Phi$ and integrity constraint $\Gamma \in \mathcal{L}_{\text{sat}}$.

Let $\boldsymbol{B}' = \boldsymbol{B} + \boldsymbol{B}$. By weak homogeneity, we have $\mathcal{F}(\boldsymbol{B},\Gamma) \subseteq \mathcal{F}(\boldsymbol{B}',\Gamma)$. Furthermore, since $n_{\boldsymbol{B}'}(\ell) = 2 \cdot n_{\boldsymbol{B}}(\ell)$ for all literals $\ell \in \Phi$, any complete and Γ-consistent judgment set that is $\succeq^{\boldsymbol{B}'}_{\text{min}}$-maximal, is also $\succeq^{\boldsymbol{B}}_{\text{min}}$-maximal. Thus, $\textsc{pr}(\boldsymbol{B}',\Gamma) \subseteq \textsc{pr}(\boldsymbol{B},\Gamma)$. Therefore, if we can show that $\mathcal{F}(\boldsymbol{B}',\Gamma) \subseteq \textsc{pr}(\boldsymbol{B}',\Gamma)$, we obtain

$$\mathcal{F}(\boldsymbol{B},\Gamma) \subseteq \mathcal{F}(\boldsymbol{B}',\Gamma) \subseteq \textsc{pr}(\boldsymbol{B}',\Gamma) \subseteq \textsc{pr}(\boldsymbol{B},\Gamma)$$

and we are done.

Suppose $N_{\boldsymbol{B}'} \geq 0$ (with $N_{\boldsymbol{B}'}$ as defined in the proof of Proposition 3.3.1). Then any $J \in \textsc{pr}(\boldsymbol{B}',\Gamma)$ contains the strict majority set M' of $\boldsymbol{B}'$ (since otherwise, by completeness, J must contain a literal ℓ with $n_{\boldsymbol{B}'}(\ell) < 0$, and $N_{\boldsymbol{B}'} < 0$). Therefore, M' is consistent with Γ and thus by weak Condorcet consistency, $M' \subseteq J$ for each $J \in \mathcal{F}(\boldsymbol{B}',\Gamma)$. But for any complete and Γ-consistent judgment set $J \in \mathcal{J}(\Gamma)$ which contains M', we have $J \in \textsc{pr}(\boldsymbol{B}',\Gamma)$ (since either $N_{\boldsymbol{B}'} > 0$ and M' itself is a complete and Γ-consistent judgment set, which is therefore the only collective decision accepted by $\textsc{pr}(\boldsymbol{B}',\Gamma)$; or $N_{\boldsymbol{B}'} = 0$ and all complete and Γ-consistent extensions J of M' have $\min_{\ell \in J} n_{\boldsymbol{B}}(\ell) = 0$, so $J \in \textsc{pr}(\boldsymbol{B}',\Gamma)$). Therefore, by collective completeness and consistency, we have $\mathcal{F}(\boldsymbol{B}',\Gamma) \subseteq \textsc{pr}(\boldsymbol{B}',\Gamma)$, which completes the proof in case $N_{\boldsymbol{B}'} \geq 0$.

For the remainder of the proof, suppose $N_{\boldsymbol{B}'} < 0$. We will show that $\mathcal{F}(\boldsymbol{B}',\Gamma) \subseteq \textsc{pr}(\boldsymbol{B}',\Gamma)$ by constructing a chain of profiles $\boldsymbol{B}^0, \ldots, \boldsymbol{B}^k$ together with corresponding strict majority profiles $\boldsymbol{B}^0(M^0), \ldots, \boldsymbol{B}^{k-1}(M^{k-1})$ (where M^i is the strict majority set of $\boldsymbol{B}^i$) such that $\boldsymbol{B}^{i+1} = \boldsymbol{B}^i + \boldsymbol{B}^i(M^i)$ for $0 \leq i < k$, and $\mathcal{F}(\boldsymbol{B}',\Gamma) \subseteq \mathcal{F}(\boldsymbol{B}^k,\Gamma)$ and $\mathcal{F}(\boldsymbol{B}^0,\Gamma) \subseteq \textsc{pr}(\boldsymbol{B}',\Gamma)$. Therefore, by weak majority profile consistency, we get the desired inclusion

$$\mathcal{F}(\boldsymbol{B}',\Gamma) \subseteq \mathcal{F}(\boldsymbol{B}^k,\Gamma) \subseteq \cdots \subseteq \mathcal{F}(\boldsymbol{B}^0,\Gamma) \subseteq \textsc{pr}(\boldsymbol{B}',\Gamma).$$

Note that since $\boldsymbol{B}' = \boldsymbol{B} + \boldsymbol{B}$, the majority margin in $\boldsymbol{B}'$ of any literal must be even. Therefore, $N_{\boldsymbol{B}'}$ is even. Let $k = -\frac{1}{2} \cdot N_{\boldsymbol{B}'}$ and consider the sets of literals with equal majority support,[2]

$$\begin{aligned}
\Lambda_1 &= \{\ell \in \Phi \mid n_{\boldsymbol{B}'}(\ell) = 2\},\\
\Lambda_2 &= \{\ell \in \Phi \mid n_{\boldsymbol{B}'}(\ell) = 4\},\\
&\vdots\\
\Lambda_k &= \{\ell \in \Phi \mid n_{\boldsymbol{B}'}(\ell) = -N_{\boldsymbol{B}'}\}.
\end{aligned}$$

Define the profile $\boldsymbol{B}^0$ as

$$\boldsymbol{B}^0 = \boldsymbol{B}' + k \cdot \sum_{\ell \in M'_{2k}} \boldsymbol{B}_{\sim\ell} + \sum_{i=1}^{k} \left(i \cdot \sum_{\ell \in \Lambda_i} \boldsymbol{B}_{\sim\ell} \right),$$

[2] Note that these sets are some of the equivalence classes $E_i \in [\![\succeq_{\boldsymbol{B}'}]\!]_\sim$ as defined in Section 3.2, but with different indices.

where M'_{2k} is the $2k$-qualified majority set of $\boldsymbol{B'}$. In other words, $\boldsymbol{B}^0$ is constructed from $\boldsymbol{B'}$ by adding k pairs of voters who only agree on rejecting ℓ for each literal $\ell \in \Phi$ that enjoys a majority margin strictly greater than $2k$ (and is thus necessarily accepted by $\textsc{pr}(\boldsymbol{B'}, \Gamma)$), and adding i pairs of voters who only agree on rejecting ℓ for each literal $\ell \in \Phi$ that enjoys a majority margin of exactly $2i$ for $1 \leq i \leq k$. Thus, the majority margins of $\boldsymbol{B'}$ and $\boldsymbol{B}^0$ are related as

$$n_{\boldsymbol{B}^0}(\ell) = \begin{cases} n_{\boldsymbol{B'}}(\ell) - 2k & \text{if } n_{\boldsymbol{B'}}(\ell) > 2k; \\ n_{\boldsymbol{B'}}(\ell) + 2k & \text{if } n_{\boldsymbol{B'}}(\ell) < 2k; \\ 0 & \text{else.} \end{cases}$$

And therefore, the strict majority set M^0 of $\boldsymbol{B}^0$ equals the $2k$-qualified majority set M'_{2k} of $\boldsymbol{B'}$.

But all literals in M'_{2k} are necessarily accepted by $\textsc{pr}(\boldsymbol{B'}, \Gamma)$ (since by definition of $N_{\boldsymbol{B'}} = -2k$, all literals $\ell \in \Phi$ with $n_{\boldsymbol{B'}}(\ell) < -2k$ are necessarily rejected by $\textsc{pr}(\boldsymbol{B'}, \Gamma)$), and thus the set $M^0 = M'_{2k}$ is consistent with Γ. Therefore, by weak Condorcet consistency, we have $M^0 \subseteq J$ for all $J \in \mathcal{F}(\boldsymbol{B}^0, \Gamma)$, and by collective completeness and consistency, $J \in \mathcal{J}(\Gamma)$ for all $J \in \mathcal{F}(\boldsymbol{B}^0, \Gamma)$. Since $\textsc{pr}(\boldsymbol{B'}, \Gamma)$ contains exactly those complete and Γ-consistent judgment sets which extend M'_{2k} (because otherwise $N_{\boldsymbol{B'}} > -2k$ by some $\succeq^{\boldsymbol{B'}}_{\min}$-maximal judgment set $J \in \mathcal{J}(\Gamma)$ with $\min_{\ell \in J} n_{\boldsymbol{B'}}(\ell) > -2k$), we obtain $\mathcal{F}(\boldsymbol{B}^0, \Gamma) \subseteq \textsc{pr}(\boldsymbol{B'}, \Gamma)$.

We now construct the profile $\boldsymbol{B}^{i+1}$ for $0 \leq i < k$ recursively from $\boldsymbol{B}^i$ by defining a strict majority profile $\boldsymbol{B}^i(M^i)$ of $\boldsymbol{B}^i$ and setting $\boldsymbol{B}^{i+1} = \boldsymbol{B}^i + \boldsymbol{B}^i(M^i)$. The profile $\boldsymbol{B}^i(M^i)$ is defined as

$$\boldsymbol{B}^i(M^i) = \sum_{\ell \in M^i \cup \Lambda_{k-i}} \boldsymbol{B}_\ell.$$

In other words, $\boldsymbol{B}^i(M^i)$ contains a pair of voters which only agree on accepting ℓ for each ℓ which receives strict majority support in $\boldsymbol{B}^i$ and for each ℓ which enjoys a majority margin of $2(k-i)$ in $\boldsymbol{B'}$. Since we define $\boldsymbol{B}^i(M^i)$ recursively, $\boldsymbol{B}^0(M^0)$ contains minimal ℓ-profiles for all $\ell \in \Lambda_k$, $\boldsymbol{B}^1(M^1)$ contains minimal ℓ-profiles for all $\ell \in \Lambda_{k-1}$, etc.

Note that we have $M^{i+1} = M^i \cup \Lambda_{k-i}$ for $0 \leq i < k$. Thus, starting from profile $\boldsymbol{B}^0$ where M^0 contains those literals $\ell \in \Phi$ with $n_{\boldsymbol{B'}}(\ell) > 2k$ and $n_{\boldsymbol{B}^0}(\ell') = 0$ for all $\ell' \in \bigcup_{i=1}^k \Lambda_i$, we generate strict majority support in $\boldsymbol{B}^1$ for the literals $\ell' \in \Lambda_k$, then generate strict majority support in $\boldsymbol{B}^2$ for the literals $\ell' \in \Lambda_{k-1}$, etc. This implies that all $\ell' \in \Lambda_{k-i}$ have $n_{\boldsymbol{B}^i}(\ell') = 0$, while all $\ell \in \Lambda_j$ for $j > k-i$ have $\ell \in M^i$. Thus, $\boldsymbol{B}^i(M^i)$ is indeed a strict majority profile of $\boldsymbol{B}^i$.

It only remains to be shown that $\mathcal{F}(\boldsymbol{B}', \Gamma) \subseteq \mathcal{F}(\boldsymbol{B}^k, \Gamma)$. Note that $\boldsymbol{B}^k$ has been recursively constructed as

$$\begin{aligned}
\boldsymbol{B}^k &= \boldsymbol{B}^0 + \boldsymbol{B}^0(M^0) + \boldsymbol{B}^1(M^1) + \cdots + \boldsymbol{B}^k(M^k) \\
&= \left(\boldsymbol{B}' + k \cdot \left(\sum_{\ell \in M'_{2k}} \boldsymbol{B}_{\sim\ell}\right) + 1 \cdot \left(\sum_{\ell \in \Lambda_1} \boldsymbol{B}_{\sim\ell}\right) + \cdots + k \cdot \left(\sum_{\ell \in \Lambda_k} \boldsymbol{B}_{\sim\ell}\right)\right) \\
&\quad + \left(\sum_{\ell \in M^0 \cup \Lambda_k} \boldsymbol{B}_\ell\right) + \left(\sum_{\ell \in M^0 \cup \Lambda_k \cup \Lambda_{k-1}} \boldsymbol{B}_\ell\right) + \cdots + \left(\sum_{\ell \in M^0 \cup \Lambda_k \cup \cdots \cup \Lambda_1} \boldsymbol{B}_\ell\right).
\end{aligned}$$

Since $M'_{2k} = M^0$, we can rearrange the terms in this sum as

$$\begin{aligned}
\boldsymbol{B}^k = \boldsymbol{B}' + k \cdot &\left(\sum_{\ell \in M'_{2k}} \boldsymbol{B}_\ell + \boldsymbol{B}_{\sim\ell}\right) \\
&+ 1 \cdot \left(\sum_{\ell \in \Lambda_1} \boldsymbol{B}_\ell + \boldsymbol{B}_{\sim\ell}\right) + \cdots + k \cdot \left(\sum_{\ell \in \Lambda_k} \boldsymbol{B}_\ell + \boldsymbol{B}_{\sim\ell}\right).
\end{aligned}$$

Since all majority margins of each term $\boldsymbol{B}_\ell + \boldsymbol{B}_{\sim\ell}$ are zero, we conclude that $\boldsymbol{B}^k = \boldsymbol{B}' + \boldsymbol{B}_{\mathrm{E}}$ for an E-profile $\boldsymbol{B}_{\mathrm{E}}$. Thus, by weak E-invariance, we have $\mathcal{F}(\boldsymbol{B}', \Gamma) \subseteq \mathcal{F}(\boldsymbol{B}^k, \Gamma)$, which completes the proof. □

Theorem 3.3.2 is the analogue of Theorem 17 of Lamboray [2009a]. Note however, that in the case of direct ordinal aggregation, the prudent rule is characterized by a slightly weaker version of weak homogeneity, which contains the extra condition that the original profile has an odd number of voters. In the case of direct democracy, we do not need to use the double profile $\boldsymbol{B}' = \boldsymbol{B} + \boldsymbol{B}$ to obtain even majority margins, if the profile has an even number of voters. This is because all majority margins in a profile where all voters vote directly, are of the same parity. Thus, profiles with an even number of voters always have even majority margins for all literals. But this is not the case if an odd number of voters delegate cyclically (or if the number of voters is odd to begin with), which is why the liquid democracy setting requires a slightly stronger axiom.

We now move on to the ranked agenda rule, and show that it satisfies all axioms of Section 2.2.3. The analogous result of Lamboray [2009b] is his Proposition 2.

Proposition 3.3.3. *The ranked agenda rule on the universal domain is collectively complete and consistent, weakly majority profile consistent, weakly qualified majority profile consistent, weakly Condorcet consistent, weakly E-invariant, weakly homogeneous, and monotonically consistent.*

Proof. Fix some profile $\boldsymbol{B} \in \mathcal{B}^\Phi$ and integrity constraint $\Gamma \in \mathcal{L}_{\mathrm{sat}}$. Due to Proposition 3.2.4, we can use the functional and algorithmic definitions of the ranked agenda rule interchangeably in this proof.

Collective completeness and consistency are immediate from the functional definition of the ranked agenda rule.

We will show that the ranked agenda rule is weakly qualified majority profile consistent, which implies that it is also weakly majority profile consistent. Let $\gamma \in \mathbb{N}_0$, let M_γ be the γ-qualified majority set of $\boldsymbol{B}$, and let $\boldsymbol{B}(M_\gamma)$ be a γ-qualified majority set of $\boldsymbol{B}$ over some set $\zeta \subseteq \{\ell \in \Phi \mid n_{\boldsymbol{B}}(\ell) = \gamma\}$.

The ranked agenda algorithm for any profile $\boldsymbol{B}' \in \mathcal{B}^\Phi$ considers all linear orders on Φ that are compatible with $\succeq_{\boldsymbol{B}'}$, and accepts literals in that order unless the literal renders the possible collective judgment set inconsistent with Γ. But by Fact 2.2.25, the majority margins of $\ell \in \Phi$ under $\boldsymbol{B}$ and $\boldsymbol{B} + \boldsymbol{B}(M_\gamma)$ are related as

$$n_{\boldsymbol{B}+\boldsymbol{B}(M_\gamma)}(\ell) = \begin{cases} n_{\boldsymbol{B}}(\ell) + 2 & \text{if } n_{\boldsymbol{B}}(\ell) > \gamma \text{ or } \ell \in \zeta; \\ n_{\boldsymbol{B}}(\ell) - 2 & \text{if } n_{\boldsymbol{B}}(\ell) < \gamma \text{ or } {\sim}\ell \in \zeta; \\ n_{\boldsymbol{B}}(\ell) & \text{else.} \end{cases}$$

Therefore, whenever $\ell \succeq_{\boldsymbol{B}+\boldsymbol{B}(M_\gamma)} \ell'$ for $\ell, \ell' \in \Phi$, we also have $\ell \succeq_{\boldsymbol{B}} \ell'$. Thus any linear order on Φ which is compatible with $\succeq_{\boldsymbol{B}+\boldsymbol{B}(M_\gamma)}$ is also compatible with $\succeq_{\boldsymbol{B}}$. Therefore, if $J \in \textsc{ra}(\boldsymbol{B} + \boldsymbol{B}(M_\gamma))$, we must have $J \in \textsc{ra}(\boldsymbol{B})$. We conclude that $\textsc{ra}(\boldsymbol{B} + \boldsymbol{B}(M_\gamma)) \subseteq \textsc{ra}(\boldsymbol{B})$.

For weak Condorcet consistency, suppose the strict majority set M of $\boldsymbol{B}$ is consistent with Γ. Then for each linear order which is consistent with $\succeq_{\boldsymbol{B}}$, all literals $\ell \in M$ are ranked before all other literals. And accepting all literals $\ell \in M$ is consistent with Γ, so the algorithm does so for each linear order it considers. Thus $M \subseteq J$ for all $J \in \textsc{ra}(\boldsymbol{B}, \Gamma)$.

Weak E-invariance and weak homogeneity are immediate from the algorithmic definition of the ranked agenda rule, when we note that the orders $\succeq_{\boldsymbol{B}}$, $\succeq_{\boldsymbol{B}+\boldsymbol{B}_{\mathrm{E}}}$ and $\succeq_{\boldsymbol{B}+\boldsymbol{B}}$ are all identical.

For monotonic consistency, let $\ell \in \Phi$ and $J \in \textsc{ra}(\boldsymbol{B}, \Gamma)$ such that $\ell \in J$. By Fact 2.2.23, the majority margins of $\ell' \in \Phi$ under $\boldsymbol{B}$ and $\boldsymbol{B} + \boldsymbol{B}_\ell$ are related as

$$n_{\boldsymbol{B}+\boldsymbol{B}_\ell}(\ell') = \begin{cases} n_{\boldsymbol{B}}(\ell') + 2 & \text{if } \ell' = \ell; \\ n_{\boldsymbol{B}}(\ell') - 2 & \text{if } \ell' = {\sim}\ell; \\ n_{\boldsymbol{B}}(\ell') & \text{else.} \end{cases}$$

Let $J' \in \mathcal{J}(\Gamma)$. Then by the functional definition of the ranked agenda rule, we have $J \succeq^{\boldsymbol{B}}_{\mathrm{disc}} J'$, i.e.,

$$\min_{\ell' \in J \setminus J'} n_{\boldsymbol{B}}(\ell') \geq \min_{\ell' \in J' \setminus J} n_{\boldsymbol{B}}(\ell').$$

If $\ell \in J'$, then $\ell, {\sim}\ell \notin J \setminus J'$ and $\ell, {\sim}\ell \notin J' \setminus J$. And if $\ell \notin J'$, then $\ell \in J \setminus J'$ and ${\sim}\ell \in J' \setminus J$. In both cases, the inequality

$$\min_{\ell' \in J \setminus J'} n_{\boldsymbol{B}+\boldsymbol{B}_\ell}(\ell') \geq \min_{\ell' \in J' \setminus J} n_{\boldsymbol{B}+\boldsymbol{B}_\ell}(\ell')$$

remains by the relation between $n_{\boldsymbol{B}}(\ell')$ and $n_{\boldsymbol{B}+\boldsymbol{B}_\ell}(\ell')$ above. So $J \succeq^{\boldsymbol{B}+\boldsymbol{B}_\ell}_{\mathrm{disc}} J'$. We conclude that J is $\succeq_{\boldsymbol{B}+\boldsymbol{B}_\ell}$-maximal, and $J \in \textsc{ra}(\boldsymbol{B} + \boldsymbol{B}_\ell, \Gamma)$. Therefore, $\{J \in \textsc{ra}(\boldsymbol{B}, \Gamma) \mid \ell \in J\} \subseteq \textsc{ra}(\boldsymbol{B} + \boldsymbol{B}_\ell, \Gamma)$. $\square$

The following theorem shows that the ranked agenda rule can be characterized as the largest rule (in the sense of set-inclusion) which satisfies the same axioms as the prudent rule, supplemented by monotonic consistency. In other words, the ranked agenda rule is the largest prudent rule which is monotone under adding minimal ℓ-profiles. It is a generalization of Theorem 2 of Lamboray [2009b].

Theorem 3.3.4. *The ranked agenda rule is the $\subseteq$-largest liquid aggregation rule on the universal domain that is collectively complete and consistent, weakly majority profile consistent, weakly Condorcet consistent, weakly E-invariant, weakly homogeneous, and monotonically consistent.*

Proof. By Proposition 3.3.3, the ranked agenda rule satisfies all the properties in the theorem. What remains to be shown is that for any liquid aggregation rule $\mathcal{F}$ on the universal domain that satisfies these properties, we have $\mathcal{F}(\boldsymbol{B},\Gamma) \subseteq \textsc{ra}(\boldsymbol{B},\Gamma)$ for all $\boldsymbol{B} \in \mathcal{B}^\Phi$ and $\Gamma \in \mathcal{L}_{\text{sat}}$. We will do so by showing that for any judgment set $J' \subseteq \Phi$, if $J' \notin \textsc{ra}(\boldsymbol{B},\Gamma)$, then $J' \notin \mathcal{F}(\boldsymbol{B},\Gamma)$. Thus, fix some profile $\boldsymbol{B} \in \mathcal{B}^\Phi$ and integrity constraint $\Gamma \in \mathcal{L}_{\text{sat}}$, and let $J' \subseteq \Phi$ such that $J' \notin \textsc{ra}(\boldsymbol{B},\Gamma)$.

If J' is not complete and consistent with Γ, then $J' \notin \mathcal{F}(\boldsymbol{B},\Gamma)$ by collective completeness and consistency. Therefore, we assume in the rest of the proof that $J' \in \mathcal{J}(\Gamma)$.

Since $J' \notin \textsc{ra}(\boldsymbol{B},\Gamma)$ while $J' \in \mathcal{J}(\Gamma)$, the functional definition of the ranked agenda rule implies that there is some judgment set $J \in \mathcal{J}(\Gamma)$ such that $J \succ^{\boldsymbol{B}}_{\text{disc}} J'$, i.e.,

$$\min_{\ell \in J \setminus J'} n_{\boldsymbol{B}}(\ell) > \min_{\ell \in J' \setminus J} n_{\boldsymbol{B}}(\ell). \tag{1}$$

We are going to construct a profile $\boldsymbol{B}'$ such that $J' \notin \textsc{pr}(\boldsymbol{B}',\Gamma)$, and that $J' \in \mathcal{F}(\boldsymbol{B},\Gamma)$ implies $J' \in \mathcal{F}(\boldsymbol{B}',\Gamma)$. Since $\mathcal{F}$ satisfies all conditions in Theorem 3.3.2, $J' \notin \textsc{pr}(\boldsymbol{B}',\Gamma)$ implies that $J' \notin \mathcal{F}(\boldsymbol{B}',\Gamma)$, and thus $J' \notin \mathcal{F}(\boldsymbol{B},\Gamma)$ as desired.

Consider the set of literals which are accepted in J and J', and enjoy a majority margin in $\boldsymbol{B}$ of at most $\min_{\ell' \in J' \setminus J} n_{\boldsymbol{B}}(\ell')$,

$$L = \left\{ \ell \in J \cap J' \,\middle|\, n_{\boldsymbol{B}}(\ell) \leq \min_{\ell' \in J' \setminus J} n_{\boldsymbol{B}}(\ell') \right\}.$$

Suppose $L = \emptyset$, and set $\boldsymbol{B}' = \boldsymbol{B}$. Then trivially, $J' \in \mathcal{F}(\boldsymbol{B},\Gamma)$ implies $J' \in \mathcal{F}(\boldsymbol{B}',\Gamma)$. To show $J' \notin \textsc{pr}(\boldsymbol{B}',\Gamma)$, let $\ell \in J$. If $\ell \in J'$, then since $L = \emptyset$, we have

$$n_{\boldsymbol{B}}(\ell) > \min_{\ell' \in J' \setminus J} n_{\boldsymbol{B}}(\ell') \geq \min_{\ell' \in J'} n_{\boldsymbol{B}}(\ell').$$

And if $\ell \notin J'$, then since $J \succ^{\boldsymbol{B}}_{\text{disc}} J'$, we have

$$n_{\boldsymbol{B}}(\ell) \geq \min_{\ell' \in J \setminus J'} n_{\boldsymbol{B}}(\ell') > \min_{\ell' \in J' \setminus J} n_{\boldsymbol{B}}(\ell') \geq \min_{\ell' \in J'} n_{\boldsymbol{B}}(\ell').$$

In both cases, we have $n_{\boldsymbol{B}}(\ell) > \min_{\ell' \in J'} n_{\boldsymbol{B}}(\ell')$, and thus

$$\min_{\ell \in J} n_{\boldsymbol{B}}(\ell) > \min_{\ell' \in J'} n_{\boldsymbol{B}}(\ell').$$

In other words, $J \succ^{\boldsymbol{B}}_{\min} J'$ and therefore $J' \notin \textsc{pr}(\boldsymbol{B}, \Gamma) = \textsc{pr}(\boldsymbol{B}', \Gamma)$.

Suppose $L \neq \emptyset$. We will construct $\boldsymbol{B}'$ as an extension of $\boldsymbol{B}$ and use monotonic consistency to show that $J' \notin \textsc{pr}(\boldsymbol{B}', \Gamma)$, and that $J' \in \mathcal{F}(\boldsymbol{B}, \Gamma)$ implies $J' \in \mathcal{F}(\boldsymbol{B}', \Gamma)$.

For each $\ell \in L$, let

$$k_\ell = \left\lceil \frac{\min_{\ell' \in J' \setminus J} n_{\boldsymbol{B}}(\ell') - n_{\boldsymbol{B}}(\ell) + 1}{2} \right\rceil, \tag{2}$$

which is a strictly positive integer since $n_{\boldsymbol{B}}(\ell) \leq \min_{\ell' \in J' \setminus J} n_{\boldsymbol{B}}(\ell')$ for $\ell \in L$. We define the profile $\boldsymbol{B}'$ as

$$\boldsymbol{B}' = \boldsymbol{B} + \sum_{\ell \in L} k_\ell \cdot \boldsymbol{B}_\ell.$$

Since $\ell \in J'$ for all $\ell \in L$, monotonic consistency applied $\sum_{\ell \in L} k_\ell$ times implies that if $J' \in \mathcal{F}(\boldsymbol{B}, \Gamma)$ then $J' \in \mathcal{F}(\boldsymbol{B}', \Gamma)$, as desired.

To finally show that $J' \notin \textsc{pr}(\boldsymbol{B}', \Gamma)$, we will show $J \succ^{\boldsymbol{B}'}_{\min} J'$. In other words, $n_{\boldsymbol{B}'}(\ell) > \min_{\ell' \in J'} n_{\boldsymbol{B}'}(\ell')$ for all $\ell \in J$. Note that the majority margins of $\ell' \in \Phi$ under $\boldsymbol{B}$ and $\boldsymbol{B}'$ are related as

$$n_{\boldsymbol{B}'}(\ell') = \begin{cases} n_{\boldsymbol{B}}(\ell') + 2k_{\ell'} & \text{if } \ell' \in L; \\ n_{\boldsymbol{B}}(\ell') - 2k_{\ell'} & \text{if } {\sim}\ell' \in L; \\ n_{\boldsymbol{B}}(\ell') & \text{else.} \end{cases} \tag{3}$$

Since for each $\ell' \in J' \setminus J$, we have $\ell' \notin J$ and ${\sim}\ell' \notin J'$, and thus $\ell', {\sim}\ell' \notin L$, we obtain

$$\min_{\ell' \in J' \setminus J} n_{\boldsymbol{B}'}(\ell') = \min_{\ell' \in J' \setminus J} n_{\boldsymbol{B}}(\ell'). \tag{4}$$

Similarly,

$$\min_{\ell' \in J \setminus J'} n_{\boldsymbol{B}'}(\ell') = \min_{\ell' \in J \setminus J'} n_{\boldsymbol{B}}(\ell'). \tag{5}$$

Let $\ell \in J$. We make a case distinction.

- Suppose $\ell \notin J'$. Then since $\ell \in J \setminus J'$, we have $n_{\boldsymbol{B}'}(\ell) \geq \min_{\ell' \in J \setminus J'} n_{\boldsymbol{B}'}(\ell')$. Therefore,

$$n_{\boldsymbol{B}'}(\ell) \geq \min_{\ell' \in J \setminus J'} n_{\boldsymbol{B}'}(\ell') \underset{(5)}{=} \min_{\ell' \in J \setminus J'} n_{\boldsymbol{B}}(\ell') \underset{(1)}{>} \min_{\ell' \in J' \setminus J} n_{\boldsymbol{B}}(\ell').$$

- Suppose $\ell \in J'$ and $\ell \notin L$. Then ${\sim}\ell \notin J'$, so ${\sim}\ell \notin L$. And by definition of L, we have $n_{\boldsymbol{B}}(\ell) > \min_{\ell' \in J' \setminus J} n_{\boldsymbol{B}}(\ell)$. Therefore,

$$n_{\boldsymbol{B}'}(\ell) \underset{(3)}{=} n_{\boldsymbol{B}}(\ell) > \min_{\ell' \in J' \setminus J} n_{\boldsymbol{B}}(\ell').$$

- Suppose $\ell \in J'$ and $\ell \in L$. Then,

$$n_{\boldsymbol{B}'}(\ell) \underset{(3)}{=} n_{\boldsymbol{B}}(\ell) + 2k_\ell \underset{(2)}{\geq} \min_{\ell' \in J' \setminus J} n_{\boldsymbol{B}}(\ell) + 2 > \min_{\ell' \in J' \setminus J} n_{\boldsymbol{B}}(\ell)$$

In all cases, we have $n_{\boldsymbol{B}'}(\ell) > \min_{\ell' \in J' \setminus J} n_{\boldsymbol{B}}(\ell)$, from which we conclude that

$$n_{\boldsymbol{B}'}(\ell) > \min_{\ell' \in J' \setminus J} n_{\boldsymbol{B}}(\ell) \underset{(4)}{=} \min_{\ell' \in J' \setminus J} n_{\boldsymbol{B}'}(\ell') \geq \min_{\ell' \in J'} n_{\boldsymbol{B}'}(\ell').$$

This completes the proof. □

Finally, we can fully characterize the ranked agenda rule by strengthening majority profile consistency to qualified majority profile consistency, generalizing Theorem 3 of Lamboray [2009b]. That is, out of all liquid aggregation rules which satisfy the axioms in Theorem 3.3.4, the ranked agenda rule is not only the largest rule, but also the only rule that satisfies qualified majority profile consistency.

Corollary 3.3.5. *The ranked agenda rule is the only liquid aggregation rule on the universal domain that is collectively complete and consistent, weakly qualified majority profile consistent, weakly Condorcet consistent, weakly E-invariant, weakly homogeneous, and monotonically consistent.*

Proof. Let $\mathcal{F}$ be a liquid aggregation rule that satisfies the conditions in the theorem and fix some profile $\boldsymbol{B} \in \mathcal{B}^\Phi$ and integrity constraint $\Gamma \in \mathcal{L}_{\text{sat}}$. By Proposition 3.3.3, the ranked agenda rule satisfies the conditions in the theorem. And since weak qualified majority profile consistency implies weak majority profile consistency, Theorem 3.3.4 states that $\mathcal{F}(\boldsymbol{B}, \Gamma) \subseteq \textsc{ra}(\boldsymbol{B}, \Gamma)$. What remains to be shown is that $\textsc{ra}(\boldsymbol{B}, \Gamma) \subseteq \mathcal{F}(\boldsymbol{B}, \Gamma)$.

Let $J \in \textsc{ra}(\boldsymbol{B}, \Gamma)$. We will construct a chain of profiles $\boldsymbol{B}^0, \ldots, \boldsymbol{B}^k$ such that $\boldsymbol{B}^0 = \boldsymbol{B}$, and $\mathcal{F}(\boldsymbol{B}^k, \Gamma) \subseteq \cdots \subseteq \mathcal{F}(\boldsymbol{B}^0, \Gamma)$, and $\textsc{ra}(\boldsymbol{B}^k, \Gamma) = \{J\}$. Since $\mathcal{F}(\boldsymbol{B}^k, \Gamma) \subseteq \textsc{ra}(\boldsymbol{B}^k, \Gamma)$ by Theorem 3.3.4, and $\mathcal{F}(\boldsymbol{B}^k, \Gamma)$ is non-empty by definition of a liquid aggregation rule, we obtain $J \in \mathcal{F}(\boldsymbol{B}^k, \Gamma) \subseteq \mathcal{F}(\boldsymbol{B}, \Gamma)$. This proves that $\textsc{ra}(\boldsymbol{B}, \Gamma) \subseteq \mathcal{F}(\boldsymbol{B}, \Gamma)$.

For any $\boldsymbol{B}' \in \mathcal{B}^\Phi$, let

$$\Omega(\boldsymbol{B}') := \textsc{ra}(\boldsymbol{B}', \Gamma) \setminus \{J\}.$$

Let $\boldsymbol{B}^0 = \boldsymbol{B}$. If $\Omega(\boldsymbol{B}^0) = \emptyset$, then $\textsc{ra}(\boldsymbol{B}^0, \Gamma) = \{J\}$, so setting $k = 0$ completes the proof. If $\Omega(\boldsymbol{B}^0) \neq \emptyset$, let $J' \in \Omega(\boldsymbol{B}^0)$. By the functional definition of the ranked agenda rule and the observation that $J, J' \in \textsc{ra}(\boldsymbol{B}^0, \Gamma)$, we must have

$$\min_{\ell' \in J \setminus J'} n_{\boldsymbol{B}^0}(\ell') = \min_{\ell' \in J' \setminus J} n_{\boldsymbol{B}^0}(\ell') = -\gamma$$

for some $\gamma \in \mathbb{Z}$.

Suppose $\gamma < 0$ and let $\ell \in J \setminus J'$ (noting that the set $J \setminus J'$ is non-empty, since J and J' are two different complete and consistent judgment sets). Since

$$\min_{\ell' \in J \setminus J'} n_{\boldsymbol{B}^0}(\ell') = -\gamma > 0,$$

we have that $n_{\boldsymbol{B}^0}(\ell) > 0$, and therefore we have $n_{\boldsymbol{B}^0}(\sim\ell) = -n_{\boldsymbol{B}^0}(\ell) < 0$. But $\ell \notin J'$ and J' is complete, and $\ell \in J$ and J is consistent, so $\sim\ell \in J' \setminus J$. Thus, $\min_{\ell' \in J' \setminus J} n_{\boldsymbol{B}^0}(\ell') < 0$ and therefore

$$\gamma = -\min_{\ell' \in J' \setminus J} n_{\boldsymbol{B}^0}(\ell') > 0,$$

a contradiction. We conclude that $\gamma \geq 0$.

We now construct profile $\boldsymbol{B}^1$ by extending profile $\boldsymbol{B}^0$, and we show that it satisfies three expressions which allow us to derive the desired chain of set-inclusions and -equalities at the end of the proof.

Let M_γ be the γ-qualified majority set of $\boldsymbol{B}^0$. Let

$$\zeta = \{\ell \in \Phi \mid n_{\boldsymbol{B}^0}(\ell) = \gamma \text{ and } \ell \in J\}$$

and define the γ-qualified majority profile $\boldsymbol{B}^0(M_\gamma)$ of $\boldsymbol{B}^0$ as

$$\boldsymbol{B}^0(M_\gamma) = \sum_{\ell \in M_\gamma \cup \zeta} \boldsymbol{B}_\ell.$$

Let $\boldsymbol{B}^1 = \boldsymbol{B}^0 + \boldsymbol{B}^0(M_\gamma)$. Then by weak qualified majority profile consistency, we have our first expression

$$\mathcal{F}(\boldsymbol{B}^1, \Gamma) \subseteq \mathcal{F}(\boldsymbol{B}^0, \Gamma). \tag{1}$$

Secondly, we will show that $J \in \textsc{ra}(\boldsymbol{B}^1, \Gamma)$. Let $\boldsymbol{B}' = \boldsymbol{B}^0 + \sum_{\ell \in M_\gamma} \boldsymbol{B}_\ell$. The majority margins of literals $\ell \in \Phi$ under profiles $\boldsymbol{B}^0$ and $\boldsymbol{B}'$ are related as

$$n_{\boldsymbol{B}'}(\ell) = \begin{cases} n_{\boldsymbol{B}^0}(\ell) + 2 & \text{if } n_{\boldsymbol{B}^0}(\ell) > \gamma; \\ n_{\boldsymbol{B}^0}(\ell) - 2 & \text{if } n_{\boldsymbol{B}^0}(\ell) < \gamma; \\ n_{\boldsymbol{B}^0}(\ell) & \text{else.} \end{cases}$$

Thus, the orders $\succeq_{\boldsymbol{B}^0}$ and $\succeq_{\boldsymbol{B}'}$ from the algorithmic definition of the ranked agenda rule are identical. Therefore, we have $\textsc{ra}(\boldsymbol{B}^0, \Gamma) = \textsc{ra}(\boldsymbol{B}', \Gamma)$ and thus $J \in \textsc{ra}(\boldsymbol{B}', \Gamma)$. Since $\zeta \subseteq J$ and $\boldsymbol{B}^1 = \boldsymbol{B}' + \sum_{\ell \in \zeta} \boldsymbol{B}_\ell$, we have our second expression

$$J \in \textsc{ra}(\boldsymbol{B}^1, \Gamma) \tag{2}$$

by $|\zeta|$ applications of monotonic consistency.

Thirdly, we show that $\textsc{ra}(\boldsymbol{B}^1, \Gamma) \subsetneq \textsc{ra}(\boldsymbol{B}^0, \Gamma)$. By weak qualified majority profile consistency, we have $\textsc{ra}(\boldsymbol{B}^1, \Gamma) \subseteq \textsc{ra}(\boldsymbol{B}^0, \Gamma)$. Thus, if we show that $J' \notin \textsc{ra}(\boldsymbol{B}^1, \Gamma)$ (and recall that $J' \in \textsc{ra}(\boldsymbol{B}^0, \Gamma)$), then we are done. To do so, we will use the functional definition of the ranked agenda rule and show that $J \succ^{\boldsymbol{B}^1}_{\text{disc}} J'$.

By Fact 2.2.25, the majority margins of $\ell \in \Phi$ under profiles $\boldsymbol{B}^0$ and $\boldsymbol{B}^1$ are related as

$$n_{\boldsymbol{B}^1}(\ell) = \begin{cases} n_{\boldsymbol{B}^0}(\ell) + 2 & \text{if } n_{\boldsymbol{B}^0}(\ell) > \gamma, \text{ or } n_{\boldsymbol{B}^0}(\ell) = \gamma \text{ and } \ell \in J; \\ n_{\boldsymbol{B}^0}(\ell) - 2 & \text{if } n_{\boldsymbol{B}^0}(\ell) < \gamma, \text{ or } n_{\boldsymbol{B}^0}(\ell) = -\gamma \text{ and } \ell \notin J; \\ n_{\boldsymbol{B}^0}(\ell) & \text{else.} \end{cases}$$

Since $\min_{\ell \in J \setminus J'} n_{\boldsymbol{B}^0}(\ell) = \min_{\ell \in J' \setminus J} n_{\boldsymbol{B}^0}(\ell) = -\gamma$, this implies that

$$\min_{\ell \in J \setminus J'} n_{\boldsymbol{B}^1}(\ell) = \min_{\ell \in J \setminus J'} n_{\boldsymbol{B}^0}(\ell) = -\gamma$$

and

$$\min_{\ell \in J' \setminus J} n_{\boldsymbol{B}^1}(\ell) = \left(\min_{\ell \in J' \setminus J} n_{\boldsymbol{B}^0}(\ell) \right) - 2 = -\gamma - 2.$$

Therefore, $J \succ_{\text{disc}}^{\boldsymbol{B}^1} J'$, which proves our third expression

$$\textsc{ra}(\boldsymbol{B}^1, \Gamma) \subsetneq \textsc{ra}(\boldsymbol{B}^0, \Gamma). \tag{3}$$

Analogously, we construct profiles $\boldsymbol{B}^2, \ldots, \boldsymbol{B}^k$ until $\Omega(\boldsymbol{B}^k) = \emptyset$, and show the equivalents of expressions (1), (2) and (3). By expressions (2) and (3), each $\textsc{ra}(\boldsymbol{B}^i)$ for $i \in \{1, \ldots, k\}$ will have strictly less elements than $\textsc{ra}(\boldsymbol{B}^{i-1})$, while still containing J. Therefore, there indeed exists a $k \in \mathbb{N}$ such that $\Omega(\boldsymbol{B}^k) = \emptyset$. And since $\Omega(\boldsymbol{B}^k) = \emptyset$, we have $\textsc{ra}(\boldsymbol{B}^k, \Gamma) = \{J\}$, as desired. By expression (1), we further have $\mathcal{F}(\boldsymbol{B}^k, \Gamma) \subseteq \cdots \subseteq \mathcal{F}(\boldsymbol{B}^0, \Gamma)$, as desired. This completes the proof. □

We conclude that if we wish to aggregate propositional judgments under arbitrary integrity constraints completely and consistently, while respecting the normative axioms of Section 2.2.3, we must use the ranked agenda rule (Corollary 3.3.5). And if we do not necessarily need to respect qualified majority profile consistency, we can only use refinements of the ranked agenda rule (Theorem 3.3.4). By Proposition 3.1.3, we further conclude that if we want to respect the axioms of Section 2.2.3, we cannot uphold resoluteness, staticity and independence, and we need a contextual understanding of neutrality and unbiasedness.

Note that all of these results hold on the universal domain, i.e., when we do not impose any individual rationality or acyclicity conditions on the profiles. Under such domain restrictions, the ranked agenda rule still satisfies the axioms by which it is characterized on the universal domain, but such domain restrictions possibly allow for other interesting liquid aggregation rules which still satisfy the axioms of Section 2.2.3.

The remainder of this thesis builds towards a refinement of the ranked agenda rule, which explicitly takes into account the delegation structure of liquid aggregation profiles.

Chapter 4

Delegation Structure and Epistemic Performance

As stated in Section 2.3, a common method of aggregating judgments in liquid democracy is constructing a proxy profile and applying a classical judgment aggregation rule to it. The ranked agenda rule is an example of this method. However, proxy profiles do not take into account the structure of the delegation graph, other than the size of its clusters. In this chapter, we examine whether the structure of a delegation graph can tell us something about the preference of voters, which is not reflected in the proxy profile. We do so by defining an information-economic model in which boundedly rational voters choose to vote directly or delegate to another voter, and numerically evaluating the correlation between (some simple measures on) the resulting delegation structure and how accurately the proxy profile reflects the voters' preferences (i.e., the epistemic performance of the profile).

In Section 4.1, we motivate our investigation of delegation structures and its main assumptions. In Section 4.2, we introduce a computational model of voter behavior in a liquid democracy setting. In Section 4.3, we describe our experiments and discuss the results. And in Section 4.4, we draw the main conclusions.

4.1 Motivation

Although judgment aggregation theorists generally take the profile of an aggregation instance as a given, it can be interesting to examine how profiles arise as a function of voter preference and behavior. Voter preferences can be dynamic, complex and hard to capture in a simple, consistent format. Therefore, any type of ballot is unavoidably just an approximation of the voter's full preference, and voters must decide how to make that approximation. Information-economic models (such as the one proposed by Green-Armytage [2015] which we described in Section 1.2.2) can shed light on part of the profile formation process.

Clearly, such models must make some simplifying assumptions about voter preference and behavior. We will use the following story as a running example to motivate our assumptions.

Example 4.1.1. In the spirit of liquid democracy's first proponent, Dodgson [1884] (alias Lewis Carroll), suppose the Queen of Wonderland organizes a liquid democracy instance, where one of the issues regards animal welfare. Alice is concerned about the well-being of animals in general, but knows very little about the policy area. However, her friend Bob has a degree in zoology and he generally shares many ideological convictions with Alice, including the value of animals. Alternatively, Alice could delegate her vote to Charlie, a veterinarian who seems to care very little about his patients, or David, who is a loving cat-owner, but knows next to nothing about other animals. Alice decides to delegate her vote to Bob, since he shares her values and has more expertise in the relevant field than she does. △

Generally, we can expect a voter to delegate to a peer if and only if the voter believes that the peer shares her values in the relevant policy area, and is a greater expert in the field. If a peer either has less (or similar) expertise than the voter, or holds a diametrically opposed opinion to her, the voter will not delegate to this peer. We will use these two criteria to model voter behavior when casting a liquid ballot.

Although these criteria seem natural, if many voters delegate to someone they closely agree with and these voters form a long enough chain, the first voter in the chain might strongly disagree with her representative at the other end of the chain. Such 'informational cascade'-like situations might negatively affect how closely the proxy profile resulting from these delegations, reflects the true opinions of the voters. Therefore, we expect that the longer the chains in a delegation structure are, the less accurate the proxy profile is with respect to voters' true opinions. We test this hypothesis by running a number of experiments on our computational model.

The results of such an analysis can be used to evaluate different liquid aggregation rules on how accurately their collective decision reflects the voters' preferences. Although from a normative, procedural standpoint (especially concerning the principle of 'one man, one vote'), we might want to base collective decisions primarily on the vote count according to a classic proxy profile, we can, at the very least, use structural considerations to break ties between liquid profiles that generate the same proxy profile.

4.2 A Model of Profile Formation

In this section, we define a model of profile formation and develop a number of measures to study the profiles generated by the model. We conclude the section with three visual examples of generated profiles to which we apply our measures.

4.2.1 Profile Formation

In our model, we consider a single issue for which a set $\mathcal{N}$ of voters have to create a delegation structure. The spectrum of possible policies concerning the issue is modeled as a one-dimensional interval $[0,1] \subseteq \mathbb{R}$. For example, position $0 \in [0,1]$ represents not addressing the issue at all, while position $1 \in [0,1]$ represents very drastic measures to address the issue. Position $0.5 \in [0,1]$ then represents a moderate approach. Each voter $i \in \mathcal{N}$ is expected to either cast a direct vote $b_i \in [0,1]$, or delegate their vote to a peer.

Each voter $i \in \mathcal{N}$ has a uniformly random position $\pi_i \in [0,1]$ in the opinion space. A voter's position in this opinion space can be interpreted as her ideological position, or the position that best serves her self-interest. However, voters can be uncertain about their exact position. In Example 4.1.1, Alice is concerned with animal well-being (hence she might know that she is 'above 0.5'), but she knows too little about animal welfare policy to pinpoint her exact ideological position on the spectrum. Therefore, each voter $i \in \mathcal{N}$ has a uniformly random 'ignorance' (i.e., the opposite of expertise) $\iota_i \in (0, \iota_{\max}]$ for some fixed $\iota_{\max} \in \mathbb{R}_{>0}$. The higher ι_i, the more uncertain voter i is about her true position: when casting a direct vote, voter i will submit a ballot $b_i \in [0,1]$ drawn from the uniform distribution on $[\pi_i - \iota_i, \pi_i + \iota_i] \cap [0,1]$. Figure 4.1 shows the position and ignorance of two voters i and j in two different initializations of the model.

In the phase of profile formation, each voter i has access to a set $\mathcal{K}_i$ of k different voters (for a fixed value $k \in \mathbb{N}$), who are drawn from $\mathcal{N} \setminus \{i\}$ uniformly and independently. Just as a voter has no access to her own exact position, she has no access to her peers' positions. However, she does observe the Euclidean distance $\delta_{i,j}$ between herself and each peer $j \in \mathcal{K}_i$, and the ignorance ι_j of each peer. In Example 4.1.1, this is reflected in Alice's observation that Bob and David (who care about animals, like Alice) are ideologically closer to her than Charlie (who cares very little about his animal patients), and that Bob and

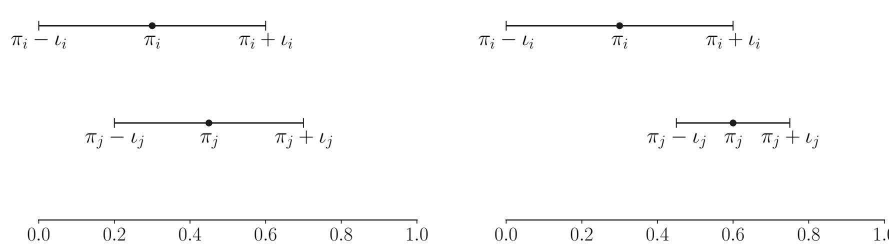

(a) Ignorance of j contains π_i. (b) Ignorance of j does not contain π_i.

Figure 4.1: Position and ignorance of two voters.

Charlie (who are experts on animal-welfare) are less ignorant than David (who knows next to nothing about animals).

Having observed the above information, each voter i must decide who to delegate to. As an example, we first consider how fully rational voters would go about this, and then present a boundedly rational version of this reasoning.

Example 4.2.1. Assuming voters aim to express their preference as accurately as possible, a fully rational voter would observe all information available to her and compute the course of action that minimizes the expectation value of the difference between her (representative's) vote and her true position, as follows.

Let $i \in \mathcal{N}$ be a rational voter. If i decides to cast a direct ballot b_i, then b_i is drawn uniformly from $[\pi_i - \iota_i, \pi_i + \iota_i]$ (assuming this interval is contained in $[0,1]$, which is the case with probability 1 in the limit $\iota_{\max} \to 0$). Since i knows the value of ι_i, she can compute that the expected error of this vote is

$$\begin{aligned}\mathbb{E}(|b_i - \pi_i|) &= \mathbb{E}(b_i - \pi_i \mid b_i \geq \pi_i) \cdot \mathbb{P}(b_i \geq \pi_i) \\ &\quad + \mathbb{E}(\pi_i - b_i \mid b_i < \pi_i) \cdot \mathbb{P}(b_i < \pi_i) \\ &= \frac{\iota_i}{2} \cdot \frac{1}{2} + \frac{\iota_i}{2} \cdot \frac{1}{2} \\ &= \frac{\iota_i}{2}, \end{aligned} \tag{1}$$

In the second equality, $\mathbb{E}(b_i - \pi_i \mid b_i \geq \pi_i) = \frac{\iota_i}{2}$ because under the assumption that $b_i \geq \pi_i$, the distribution of b_i reduces to the uniform distribution on $[\pi_i, \pi_i + \iota_i]$, of which the expectation value is $\pi_i + \frac{\iota_i}{2}$; and similarly for $\mathbb{E}(\pi_i - b_i \mid b_i < \pi_i)$.

Voter i now considers delegating to peer $j \in \mathcal{K}_i$. Assume without loss of generality that $\pi_j \geq \pi_i$. We distinguish two cases.

- If $\iota_j \geq \delta_{i,j}$ (i.e., the ignorance range of j contains π_i, see Figure 4.1a), then i can compute that the expected error (with respect to π_i) of a direct vote by j is

$$\begin{aligned}\mathbb{E}(|b_j - \pi_i|) &= \mathbb{E}(b_j - \pi_i \mid b_j \geq \pi_i) \cdot \mathbb{P}(b_j \geq \pi_i) \\ &\quad + \mathbb{E}(\pi_i - b_j \mid b_j < \pi_i) \cdot \mathbb{P}(b_j < \pi_i) \\ &= \frac{\iota_j + \delta_{i,j}}{2} \cdot \frac{\iota_j + \delta_{i,j}}{2\iota_j} + \frac{\iota_j - \delta_{i,j}}{2} \cdot \frac{\iota_j - \delta_{i,j}}{2\iota_j} \\ &= \frac{(\iota_j + \delta_{i,j})^2}{4\iota_j} + \frac{(\iota_j - \delta_{i,j})^2}{4\iota_j} \\ &= \frac{{\iota_j}^2 + {\delta_{i,j}}^2}{2\iota_j}. \end{aligned}$$

  In the second equality, $\mathbb{E}(b_j - \pi_i \mid b_j \geq \pi_i) = \frac{\iota_i + \delta_{i,j}}{2}$ because under the assumptions that $b_j \geq \pi_i$ and $\iota_j \geq \delta_{i,j}$, the distribution of b_j reduces to the uniform distribution on $[\pi_i, \pi_j + \iota_j] = [\pi_i, \pi_j] \cup [\pi_j, \pi_j + \iota_j]$. Thus $b_j - \pi_i$ is uniformly distributed on $[0, \delta_{i,j} + \iota_j]$ and its expectation value is $\frac{\delta_{i,j} + \iota_j}{2}$. Similarly, $\mathbb{P}(b_j \geq \pi_i) = \frac{\delta_{i,j} + \iota_j}{2\iota_j}$, since the interval on which

$b_j \geq \pi_i$ is $[\pi_i, \pi_j + \iota_j] = [\pi_i, \pi_j] \cup [\pi_j, \pi_j + \iota_j]$, which has length $\delta_{i,j} + \iota_j$. This is a fraction $\frac{\delta_{i,j}+\iota_j}{2\iota_j}$ of the total length of the distribution domain $[\pi_j - \iota_j, \pi_j + \iota_j]$ of b_j. The expectation value and probability for the case $b_j < \pi_i$ is analogous with interval $[\pi_j - \iota_j, \pi_i]$ of length $\iota_j - \delta_{i,j}$.

- If $\iota_j < \delta_{i,j}$ (i.e., the ignorance range of j does not contain π_i, see Figure 4.1b), then $b_j - \pi_i > 0$ for all possible ballots b_j and thus i computes the expected error as

$$\mathbb{E}(|b_j - \pi_i|) = \mathbb{E}(b_j - \pi_i) = \delta_{i,j}.$$

Voter i can now conclude that the expected error of a direct vote by j (with respect to her own position) is

$$\mathbb{E}(|b_j - \pi_i|) = \begin{cases} \frac{{\iota_j}^2 + {\delta_{i,j}}^2}{2\iota_j} & \text{if } \iota_j \geq \delta_{i,j}; \\ \delta_{i,j} & \text{if } \iota_j < \delta_{i,j}. \end{cases} \tag{2}$$

A rational voter i would only delegate to a peer j, if this reduces the expected error, i.e., $\mathbb{E}(|b_j - \pi_i|) \leq \mathbb{E}(|b_i - \pi_i|)$. Assuming all voters follow this rule, further delegations by j do not (in expectation) affect the error of the final ballot with respect to π_i, since these delegations are symmetric in π_j (in expectation). △

Note that the behavior of a rational voter is in line with Example 4.1.1: one should delegate if the other voter's ignorance and ideological distance are small. However, this level of rationality is a rather strong assumption. Therefore, we model voter behavior as a probabilistic and biased approximation of the reasoning in Example 4.2.1.

When voter i considers delegating her vote, we compute the expected error of doing so for each possible delegate $j \in \mathcal{K}_i$, with the addition of a bias. This *ignorance bias* $\beta_\iota \in [0, 2] \subseteq \mathbb{R}$ expresses to what extent voters overestimate the importance of expertise for deciding on the issue at hand. That is, we introduce β_ι as a weight in the calculation of the expected error of delegating, which increases the role of ι_j and decreases the role of $\delta_{i,j}$. The latter is represented by the *distance bias* β_δ, which is set to $\beta_\delta = 2 - \beta_\iota$ to preserve continuity of the function at the point $\iota_j = \delta_{i,j}$. All in all, the biased expected error of delegating to peer j is

$$\text{error}_{i,j} = \begin{cases} \frac{\beta_\iota \cdot {\iota_j}^2 + \beta_\delta \cdot {\delta_{i,j}}^2}{2\iota_j} & \text{if } \iota_j \geq \delta_{i,j}; \\ \delta_{i,j} & \text{if } \iota_j < \delta_{i,j}. \end{cases}$$

Note that when $\beta_\iota = \beta_\delta = 1$, this error assumes the rationally correct values from expression (2). The voter now chooses a favorite delegate probabilistically, where the probability of delegating to a peer j is inversely proportional to $\text{error}_{i,j}$.

After choosing a favorite delegate j^*, the voter decides whether to delegate to this peer or to vote directly. We introduce a bias and probabilistic character to this decision as well. Let $\beta_v \in [0, 2] \subseteq \mathbb{R}$ be a *voting bias*, which expresses an overconfidence in a voter's own direct vote. That is, we introduce β_v as a weight in the calculation of the error of voting directly and approximate this error as

$$\text{error}_i = \frac{\iota_i}{2\beta_v}.$$

When $\beta_v = 1$, this error assumes the rationally correct value from expression (1). The voter finally chooses whether to vote directly or to delegate to j^* probabilistically, where the probability of voting directly is inversely proportional to error_i, and the probability of delegating is inversely proportional to error_{i,j^*}. As noted before, if the voter decides to vote directly, her ballot is drawn uniformly at random from $[\pi_i - \iota_i, \pi_i + \iota_i] \cap [0, 1]$.

Note that (in case there is no bias), a voter is unlikely to delegate to peers who are more ignorant than herself, which helps to prevent delegation cycles. However, to guarantee acyclic profiles, we only allow delegation to peers who are strictly less ignorant than the voter, i.e., we set $\text{error}_{i,j}$ to infinity if $\iota_j \geq \iota_i$. This ensures that each edge in the delegation graph corresponds to a strict decrease in ignorance, making cycles impossible. Imposing this restriction makes it easier to compare different delegation profiles to each other without having to resolve cyclicity.

This concludes the formation of the delegation structure: all voters have decided whether to delegate (and to whom) or vote directly (and with what ballot $b_i \in [0, 1]$). Furthermore, the delegation structure is guaranteed to be acyclic. The five model parameters which can be exogenously fixed are the number of voters n, the maximal ignorance $\iota_{\max}$, the number k of peers each voter considers, the ignorance bias β_ι and the voting bias β_v. We can view these parameters as characterizing the context of the liquid democracy instance: some types of issues or external factors will, e.g., increase the bias to vote directly, and others will, e.g., decrease the ignorance voters have about the topic. Depending on this context, different delegation structures might arise, which represent voters' true positions with varying accuracy, which we refer to as the epistemic performance of the liquid profile.

4.2.2 Epistemic Performance

In the remainder of this chapter, we consider profiles that are generated by the model above. That is, each profile $\boldsymbol{B} = (B_1, \ldots, B_n)$ concerns only a single issue, and the ballots B_i either have a real value $b_i \in [0, 1]$ (in which case we write $B_i = b_i$) or refer to the delegate $j \in \mathcal{N}$ of voter i (in which case we write $B_i = j$). Aggregation rules are maps $\mathcal{F}$ from the set of possible profiles to the policy spectrum $[0, 1]$. And $\text{rep}_{\boldsymbol{B}}(i)$ and $\text{path}_{\boldsymbol{B}}(i)$ respectively denote the representative of voter i and the set of delegates connecting i to $\text{rep}_{\boldsymbol{B}}(i)$ (including i and $\text{rep}_{\boldsymbol{B}}(i)$).

When we generate a liquid profile using our model, we do not only have access to all voters' ballots, but also to their 'true preferences' π_i. This allows us to compare the preference of each voter with their representative's vote in the liquid profile, which is a measure of how accurately the proxy profile reflects the voters' preferences. In line with the terminology of Green-Armytage [2015], we refer to this difference per voter as the *expressive loss*. Consequently, the *average expressive loss* over all voters is defined as

$$\textsc{ExpressiveLoss} = \frac{1}{n} \cdot \sum_{i \in \mathcal{N}} |\pi_i - b_{\mathrm{rep}_{\boldsymbol{B}}(i)}|.$$

Furthermore, we define *systematic loss* as the difference between the collective decision a proxy profile generates and the collective decision that would be achieved under full information, i.e., if $\iota_i = 0$ for all $i \in \mathcal{N}$. We use the average ballot value in the generated proxy profile as the collective decision of a liquid profile, which is defined as

$$\mathcal{F}_{\mathrm{avg}}(\boldsymbol{B}) = \frac{1}{n} \cdot \sum_{i \in \mathcal{N}} b_{\mathrm{rep}_{\boldsymbol{B}}(i)}.$$

Note that this is an explicit choice for a specific aggregation rule, for which there are clearly many alternatives, including the common approach of taking the median of the proxy profile.

The systematic loss is then defined as

$$\textsc{SystematicLoss} = \left| \mathcal{F}_{\mathrm{avg}}(\boldsymbol{B}) - \frac{1}{n} \cdot \sum_{i \in \mathcal{N}} \pi_i \right|.$$

Finally, we can compute a utility score for each voter, expressing their satisfaction with the collective decision. We define the utility score of a voter to be the difference between her true position and the collective decision. The average utility score is referred to as *social welfare*, and is defined as

$$\textsc{SocialWelfare} = \frac{1}{n} \cdot \sum_{i \in \mathcal{N}} |\pi_i - \mathcal{F}_{\mathrm{avg}}(\boldsymbol{B})|.$$

4.2.3 Delegation Structure

In order to relate the epistemic performance of a profile to its underlying delegation structure, we define two simple metrics on profiles which express how 'deep' its delegation structure is. Intuitively, we expect that deeper delegation structures give rise to more frequent or severe informational cascades.

Viscous Delegation Depth

The first measure is similar to the proposal of 'viscous democracy' by Boldi et al. [2011]: instead of fully propagating each voter's voting weight to her delegate, we multiply a voter's voting weight with a factor $\alpha \in [0, 1] \subseteq \mathbb{R}$ before propagating it to her delegate. We use the negative of the sum of resulting voting weights of all casting voters to calculate our measure of delegation depth. That is, the more voters delegate in long chains, the lower the total voting weight for the casting voters and the higher the delegation depth.

Formally, the voting weight w_i^α of a voter $i \in \mathcal{N}$ for viscosity factor $\alpha \in [0,1]$ is defined recursively as[1]

$$w_i^\alpha = 1 + \sum_{\{j \in \mathcal{N} \mid B_j = i\}} \alpha \cdot w_j^\alpha,$$

where the sum over an empty set (i.e., for the base case where nobody delegates to voter i) is defined to equal zero. Let $R = \{i \in \mathcal{N} \mid B_i = b_i \in [0,1]\}$ be the set of all casting voters. The α-viscous delegation depth is defined as

$$D_\alpha = 1 - \frac{1}{n} \cdot \sum_{i \in R} w_i^\alpha.$$

Note that if $\alpha = 0$, the sum is equal to the number of casting voters, and thus the 0-viscous delegation depth is equal to the fraction of delegating voters over all voters. And if $\alpha = 1$, the sum is equal to the total number of voters, and thus the 1-viscous delegation depth is always equal to zero.

Further note that the sum is maximal when all voters vote directly. The sum is then equal to n. And the sum is minimal when all voters form a single chain. The sum is then equal to

$$\sum_{k=0}^{n-1} \alpha^k = \frac{1-\alpha^n}{1-\alpha} \in (1, n)$$

if $\alpha \in (0,1)$ (and it is equal to 1 if $\alpha = 0$ and to n if $\alpha = 1$). Thus, by dividing the sum of the casting voters' weights by the total number of voters and subtracting it from 1, we obtain a delegation depth between 0 and 1, which is minimal if all voters vote directly and is maximal if all voters form a single chain.

Additive Delegation Depth

The second measure is referred to as the *additive delegation depth* and counts the number of times a vote is delegated when forming the proxy profile. That is, if a vote is handed down in a chain of k voters, this adds $k-1$ to the additive delegation depth for all $k-1$ transfers of the vote. In other words, the total contribution of a chain of k votes to the additive delegation depth is

$$(k-1) + (k-2) + \cdots + 0 = \frac{k^2 - k}{2}.$$

Formally, the additive delegation depth is defined as

$$D_{\text{add}} = \frac{1}{n} \cdot \sum_{i \in \mathcal{N}} (|\operatorname{path}_{\boldsymbol{B}}(i)| - 1),$$

[1]Note that the voting weight is only well-defined for acyclic profiles.

where the fraction serves as a scaling factor which keeps the delegation depth below n. Alternatively, we can interpret the additive delegation depth as the average length of the delegation paths of all voters.

Note that the additive delegation depth (just like the α-viscous delegation depth) is minimal if all voters vote directly and maximal if all voters form a singe chain. Its minimal value is 0 and its maximal value is $\frac{n-1}{2}$.

4.2.4 Examples of Generated Profiles

In Figure 4.2, we see three randomly generated profiles using the model defined in Section 4.2.1.

In the first profile, most voters decided to vote directly as a consequence of relatively low ignorance values. Thus the α-viscous delegation depth is low at

$$D_\alpha = 1 - \frac{1}{6} \cdot (1 + 1 + 1 + 1 + (1 + \alpha)) = \frac{1 - \alpha}{6} = \frac{1}{12}$$

for $\alpha = \frac{1}{2}$, and the additive delegation depth is also low at

$$D_{\text{add}} = \frac{1}{6} \cdot (0 + 0 + 1 + 0 + 0 + 0) = \frac{1}{6}.$$

Since the final vote of all voters is relatively close to their true position, the expressive loss and systematic loss are small. However, voters are located rather far from the center of the interval, leading to a low social welfare.

In the second profile, all voters delegate (directly or indirectly) to the same voter. This can be explained by observing that this voter has a very low ignorance value and is located in the middle of all voters. The delegation depths are relatively high at

$$D_\alpha = 1 - \frac{1}{6} \cdot (1 + \alpha + \alpha + \alpha^2 + \alpha^2 + \alpha) = \frac{1}{2}$$

for $\alpha = \frac{1}{2}$, and

$$D_{\text{add}} = \frac{1}{6} \cdot (1 + 2 + 2 + 1 + 1 + 0) = \frac{7}{6}.$$

The expressive loss of the right most voter is relatively large and thus the average expressive loss is moderately large, but no long informational cascades are likely to occur for such small sets of voters.

All voters in the third profile made fully rational decisions, exhibiting no biases or probabilistic behavior. Indeed, we see that voters only delegate to peers who are ideologically very close, and voters with large ignorance values do not vote directly.

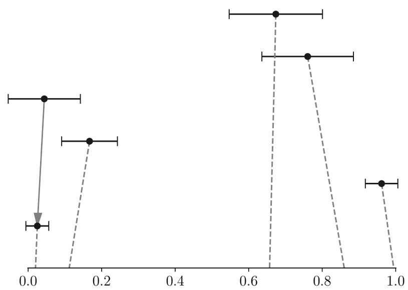

(a) Few delegations.

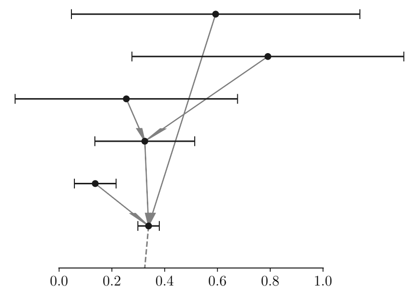

(b) Many delegations.

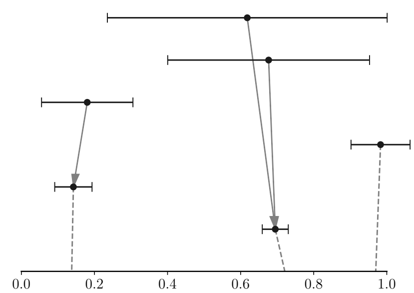

(c) Only rational delegations.

Figure 4.2: Three randomly generated liquid profiles. Voters are horizontally located at their positions on the policy spectrum, and their ignorance ranges are represented as error bars. Voters are ordered from top to bottom by decreasing ignorance. Arrows represent delegations and dotted lines represent direct votes.

4.3 Experiments and Results

The elements of Section 4.2 can be combined to conduct a number of experiments. Firstly, we study which model parameters most affect the delegation depth and epistemic performance of the profiles generated. Such an analysis is called a sensitivity analysis, and can consist of various components. We will conduct a local One-Factor-At-a-Time (OFAT) sensitivity analysis, which studies the effect of each separate model parameter on the model output while keeping the other parameters fixed, and a global Sobol' sensitivity analysis, which takes into account the interactions between parameters and can be used to discover redundant parameters.

Secondly, we use the model to study whether the delegation depth of profiles that are generated in different contexts (i.e., with different model parameters), correlates with the epistemic performance of the profiles. Moreover, we study which measure of delegation depth shows the strongest degree of correlation.

The results of these experiments can be used as an argument (not) to break ties according to (a particular measure of) delegation depth.

Before we present the experiments and results, we introduce the method of OFAT and Sobol' sensitivity analyses. We refer to the work of Ten Broeke et al. [2016] for a more thorough introduction to sensitivity analysis for agent-based models.

4.3.1 The Methods of Sensitivity Analysis

The aim of a sensitivity analysis is to map out the effect of model parameters on the model output. Let $M(\boldsymbol{x})$ be a model which operates on a vector of d inputs $\boldsymbol{x} = (x_1, \ldots, x_d) \in \mathbb{R}^d$ and returns a single output $y \in \mathbb{R}$.

One-Factor-At-a-Time (OFAT)

A One-Factor-At-a-Time (OFAT) analysis studies the local sensitivity of the model. That is, it fixes the values of all but one of the model inputs to some default value, and studies the effect of varying the other model input on the model output. The algorithm proceeds as follows.

Let i be the index of the input variable we study, let $I_i \subseteq \mathbb{R}$ be the (closed) interval of possible values of input variable i, let $\boldsymbol{x}_{-i} \in \mathbb{R}^{d-1}$ be the default values of the other input variables and let $n, m \in \mathbb{N}$. Let $x_i^1, \ldots, x_i^n \in I_i$ be equally spaced real numbers such that $x_i^1 = \min(I_i)$ and $x_i^n = \max(I_i)$. For each $j \in \{1, \ldots, n\}$, evaluate $M(\boldsymbol{x}_{-i}, x_i^j)$ in m independent runs of the model and let y_i^j be the median output value of these m model runs.

By plotting $\{y_i^j\}_{j=1}^n$ against $\{x_i^j\}_{j=1}^n$ for all $i \in \{1, \ldots, d\}$, we obtain d plots, each of which shows the effect of a single input variable on the output value of the model. Besides the size of the effect, this approach also visualizes the output as a real-valued function of each single input variable. This makes OFAT a useful method to better understand the relationship between inputs and outputs of a model. Moreover, OFAT analyses are computationally cheap compared to other methods of sensitivity analysis.

Sobol' Sensitivity Analysis

Although OFAT is perhaps the most common method to analyze complex models, it cannot detect interactions between input variables, and might depend strongly on the default values of the input variables. A Sobol' sensitivity analysis studies the global sensitivity of the model, which includes all combinations of possible model inputs. The method is due to Sobol' [2001].

The method decomposes the variance in the output of the model into variance components caused by each subset of the input variables. For the singleton sets, this translates to d different *first-order sensitivity indices* $S_1, \ldots, S_d$, defined as

$$S_i = \frac{\mathrm{Var}_{x_i}(\mathbb{E}_{\boldsymbol{x}_{-i}}(M(\boldsymbol{x}) \mid x_i))}{\mathrm{Var}_{\boldsymbol{x}}(M(\boldsymbol{x}))}.$$

In other words, S_i is the fraction of the total variance in the output of M, which can be attributed to the variance of the model output (averaged over all possible input values other than i) caused by input variable i.

Similar definitions exist for second-order effects, third-order effects, etc. Together, these orders explain the relevance of each input variable in combination with the other input variables. However, the total number of Sobol' indices grows exponentially with the number of input variables, since each subset of input variables generates a new Sobol' index. Thus, an alternative measure of the total contribution of each input variable, including interaction effects, is defined as the *total-order sensitivity indices* $S_{T_1}, \ldots, S_{T_d}$, where

$$S_{T_i} = 1 - \frac{\mathrm{Var}_{\boldsymbol{x}_{-i}}(\mathbb{E}_{x_i}(M(\boldsymbol{x}) \mid \boldsymbol{x}_{-i}))}{\mathrm{Var}_{\boldsymbol{x}}(M(\boldsymbol{x}))}.$$

In other words, S_{T_i} is the fraction of the total variance of the model, which *cannot* be attributed to the variance of the model output (averaged over all possible values of x_i) caused by all other input variables than input variable i.

To approximate the Sobol' sensitivity indices in practice, we sample the input space of the model randomly using a Saltelli sequence. This quasi-random method of sampling is due to Saltelli [2002] and generates low-discrepancy samples. Intuitively, a low-discrepancy sample is a set of points which are distributed roughly equally across space, i.e., each portion of the space with equal volume contains a roughly equal number of points. By taking a low-discrepancy sample instead of a uniformly random sample, we can reduce the number of points we sample, while still obtaining a representative set of points in the full input space.

We take two such independent samples of $N \in \mathbb{N}$ points in the d-dimensional input space, and regard them as two $N \times d$-matrices A and B. We build d more samples of N points in the d-dimensional space by defining $A_B^{(i)}$ for $i \in \{1, \ldots, d\}$ to be a copy of matrix A, where column i is replaced by column i of matrix B. The $d + 2$ matrices A, B and all $A_B^{(i)}$ now specify $N \cdot (d + 2)$ unique points in the d-dimensional input space. Using Monte Carlo estimators (i.e., a particular kind of probabilistic approximations) of variances and expectation values, we estimate the values of all S_i and S_{T_i} on these $N \cdot (d+2)$ points. These estimators converge to the real value of the Sobol' indices as we increase N, but only if N is always a power of 2 (which is why a Saltelli sample usually has size $N = 2^n$ for some $n \in \mathbb{N}$).

By plotting the approximations and confidence intervals of S_i and S_{T_i}, we can draw the following conclusions. If S_{T_i} is close to zero, then fixing the value of input parameter i does not influence the variance of the model output by much, i.e., input variable i is redundant. If S_i is close to zero, while S_{T_i} is not, then input variable i does affect the output, but mostly by higher order interactions, i.e., variable i on its own has a small effect on the output of the model. Variables with small S_i are expected to show only small effects in the OFAT analysis as well.

4.3.2 Sensitivity of Epistemic Performance

Figure 4.3 shows the OFAT analysis of the epistemic performance of the profiles generated by our model. We see that the expressive loss does not significantly depend on the total number of voters n. The expressive loss depends weakly on the number k of delegation candidates each voter considers and the ignorance bias β_ι, and it depends strongly on the maximal ignorance $\iota_{\max}$ and the voting bias β_v.

The dependence on k can be explained by observing that voters who observe more peers, have more possibilities to wisely delegate their vote and thus reduce their expressive loss. The dependence on $\iota_{\max}$ is not surprising either, as more ignorant voters are clearly less able to express their opinion precisely.

It is surprising that the expressive loss as a function of β_ι and β_v is not minimal for $\beta_\iota = 1$ and $\beta_v = 1$, since we would expect that a voter is successful at reducing her expressive loss if she has no irrational biases. However, in Figure 4.4 we see the same plot of expressive loss against β_ι and β_v, but in this model the voters behave fully rationally instead of probabilistically. In this case, we see that the expressive loss is minimal for voters without bias, i.e., with $\beta_\iota \approx \beta_v \approx 1$. We conclude that the probabilistic behavior of voters benefits those voters who either are biased towards delegating to ideologically close peers while ignoring their level of expertise, or are biased not to delegate at all.

The results of the OFAT analysis for systematic loss and social welfare show large statistical errors compared to the effect of each input variable on the output, and thus we can say less about the role of each separate input variable. However, we do see that systematic loss is small for large n, small $\iota_{\max}$ and large β_v. Clearly, if voters are less ignorant, their votes will correspond more closely to their true positions, which reduces systematic loss. And as a consequence of the famous jury theorem of Condorcet [1785], if the number of voters is large and/or many voters vote directly, their individual errors average out and the systematic error is low.

Figure 4.5 paints a similar picture of the Sobol' analysis. The first order effects of $\iota_{\max}$ and β_v on expressive loss (and to a lesser extent the first order effects of n and β_v on systematic loss) are visible, but all other first order indices are close to zero. However, we do note that the total order index of all input variables with respect to all output variables are non-zero. That is, all input variables contribute to some higher order interactions that affect the output of the model.

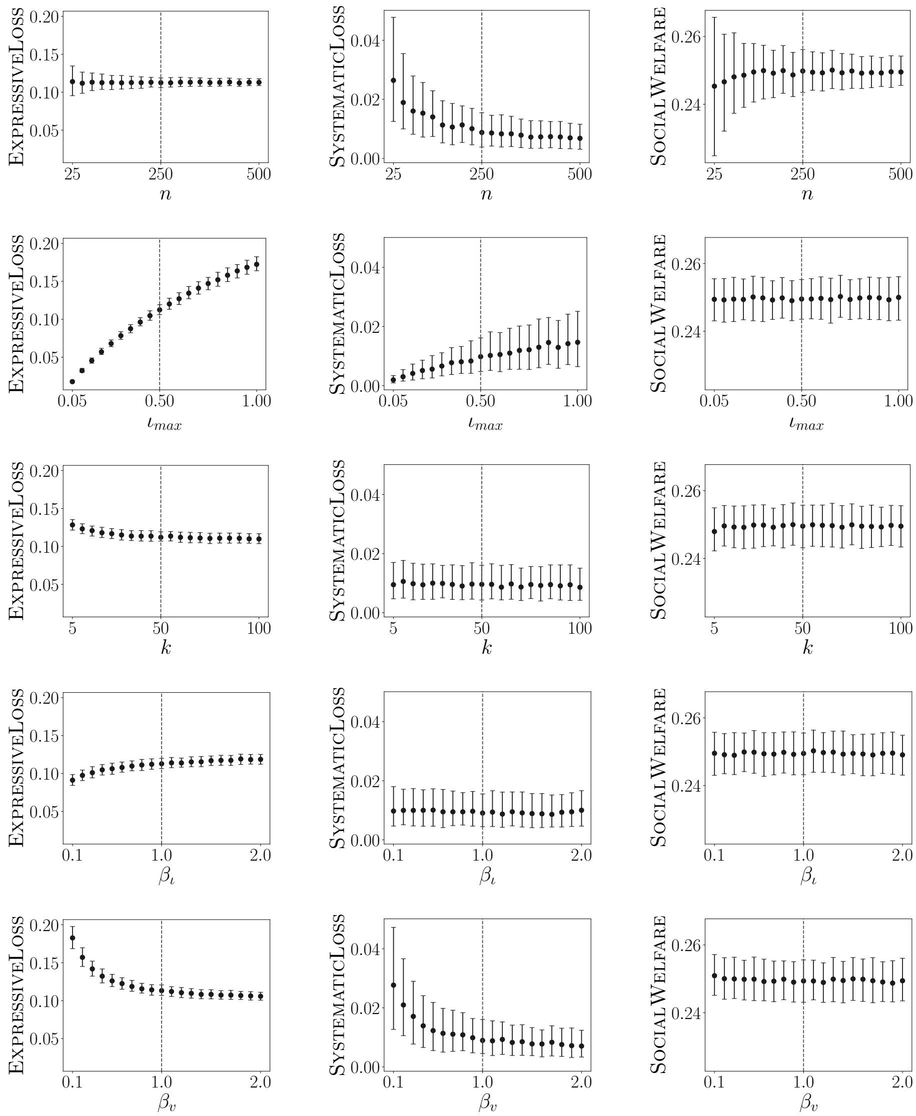

Figure 4.3: OFAT analysis of the effect of the model parameters n, $\iota_{\max}$, k, β_ι and β_v on the epistemic performance of the profiles generated by the model, where epistemic performance is measured as expressive loss, systematic loss and social welfare. Each data point is the median output over 1 000 runs. Error bars represent the first and third quartile. Dotted lines indicate the default value of each parameter.

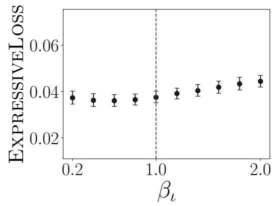

Figure 4.4: OFAT analysis of the effect of the model parameters β_ι and β_v on the expressive loss of the profiles generated by the model, where voters behave non-probabilistically. Each data point is the median output over 1 000 runs. Error bars represent the first and third quartile. Dotted lines indicate the default value of each parameter.

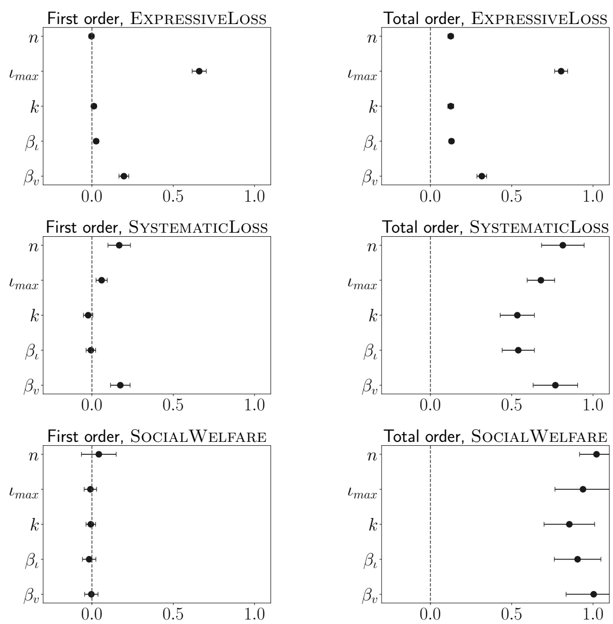

Figure 4.5: First and total order Sobol' indices for the effect of the model parameters n, $\iota_{\max}$, k, β_ι and β_v on the epistemic performance of the profiles generated by the model, where epistemic performance is measured as expressive loss, systematic loss and social welfare. The analysis was run on a Saltelli sample with $N = 4\,096$ for $d = 5$ model parameters, i.e., the model was run on $N \cdot (d + 2) = 28\,672$ samples from the input space. Error bars represent 95%-confidence intervals.

4.3.3 Sensitivity of Delegation Depth

For the sensitivity analysis of delegation depth, we applied four measures of delegation depth to the output of the profile formation model: 0-viscous delegation depth (which is equal to the fraction of delegating voters), $\frac{1}{2}$-viscous delegation depth, $\frac{9}{10}$-viscous delegation depth and additive delegation depth. We see in Figure 4.6 that our OFAT analysis showed virtually identical behavior by all these measures of delegation depth.

None of the measures of delegation depth seem to depend strongly on the total number of voters. For viscous delegation depth, this is not surprising since it is normalized by the value of n. However, recall that the additive delegation depth ranges from 0 to $\frac{n-1}{2}$ and thus it is somewhat surprising that increasing the value of n does not increase additive delegation depth. Apparently, the average length of a delegation path does not grow strongly as the number of voters grows. For small values of n, there does seem to be a small increase in delegation depth, but this increase is not statistically significant for our sample size. A possible explanation why this dependence of the additive delegation depth on the value of n is so weak, could be that the expected length of a delegation chain can be approximated by the geometric distribution,[2] if the probability that a voter delegates, is independent from the probability that a voter receives a delegation. Although this approximation is not exact since voters who receive a delegation are likely to have a small ignorance and are therefore likely not to delegate themselves, the geometric distribution predicts that for large values of n, if voters delegate with probability $P \approx 0.45$ (which is the average fraction of voters who delegate their vote, according to the plot of 0-viscous delegation depth against n in Figure 4.6), then the average length of a delegation chain is $\frac{P}{1-P} \approx 0.82$. This value is close to (but slightly higher than) the average additive delegation depth in the model. We conclude that for large values of n, the geometric approximation possibly captures the dynamics of average delegation depth as a function of n. Small values of n prohibit long delegation chains, and thus the geometric approximation does not hold, resulting in slightly lower additive delegation depths, as we indeed observed in the figure.

We further see that a larger maximal ignorance results in deeper delegation structures. This is probably due to the fact that as many voters have high expertise, it is more rational to vote directly than to delegate to another expert, who is ideologically relatively far away. Similarly, a larger number of delegation candidates per voter allows voters to find a closer ideological match to delegate to, resulting in deeper delegation structures. Finally, a bias towards voting directly intuitively decreases the delegation depth, since fewer voters decide to delegate. We also see this in the data.

A result which is not as easily understood, is that a bias towards delegating to experts (instead of to ideologically close peers) results in less deep delegation structures. Possibly, if voters weigh a possible delegate's expertise heavily, many

[2] The geometric distribution describes the number of 'tails' in a repeated coin toss (with a probability P of observing 'tails') before observing the first 'heads'. In our case, 'tails' corresponds to delegation and 'heads' to casting a direct vote.

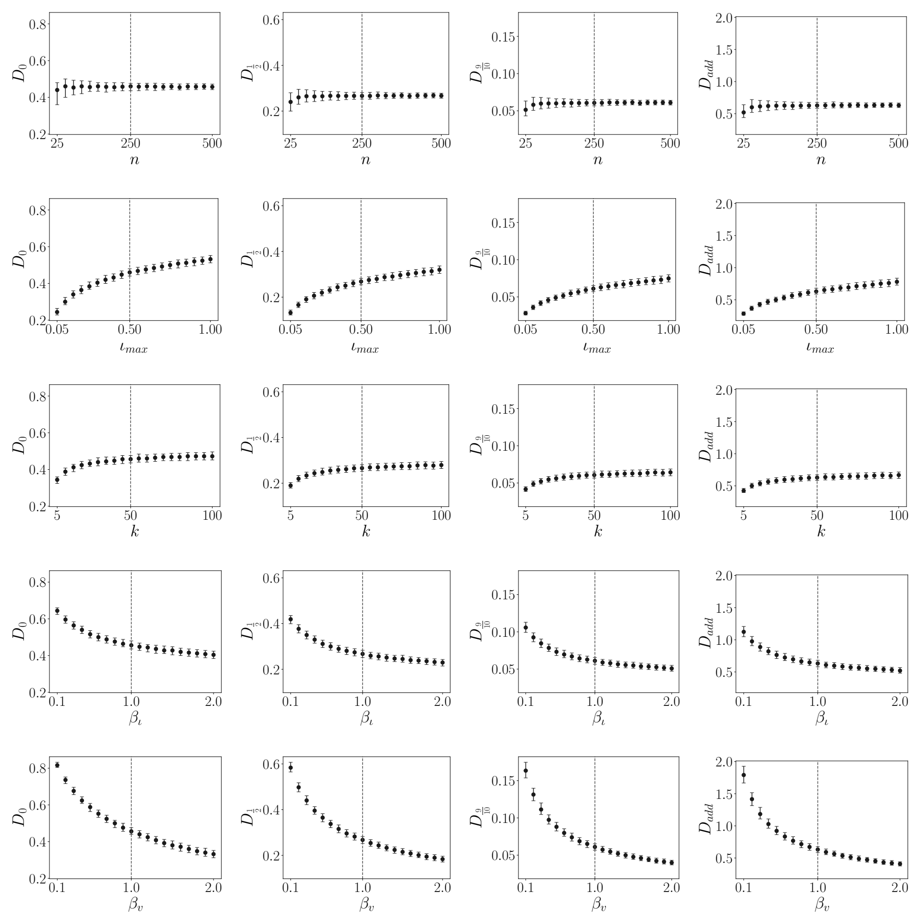

Figure 4.6: OFAT analysis of the effect of the model parameters n, $\iota_{\max}$, k, β_ι and β_v on the depth of the profiles generated by the model, where depth is measured as 0-viscous (i.e., the fraction of delegating voters), $\frac{1}{2}$-viscous, $\frac{9}{10}$-viscous and additive delegation depth. Each data point is the median output over 1 000 runs. Error bars represent the first and third quartile. Dotted lines indicate the default value of each parameter.

voters delegate directly to a clear expert instead of delegating to a non-expert who is ideologically close and who, in turn, re-delegates the vote to the clear expert.

The Sobol' analysis in Figure 4.7 paints a similar picture to the OFAT analysis: the influence of the total number of voters on delegation depth is negligible, and the voting bias and maximal ignorance have the strongest influence on delegation depth. A notable difference between the Sobol' analysis of epistemic performance and of delegation depth is that the total order Sobol' index of delegation depth is not much larger than its first order Sobol' index. That is, higher order interactions between model parameters do not play a large role in determining the delegation depth of the profiles generated.

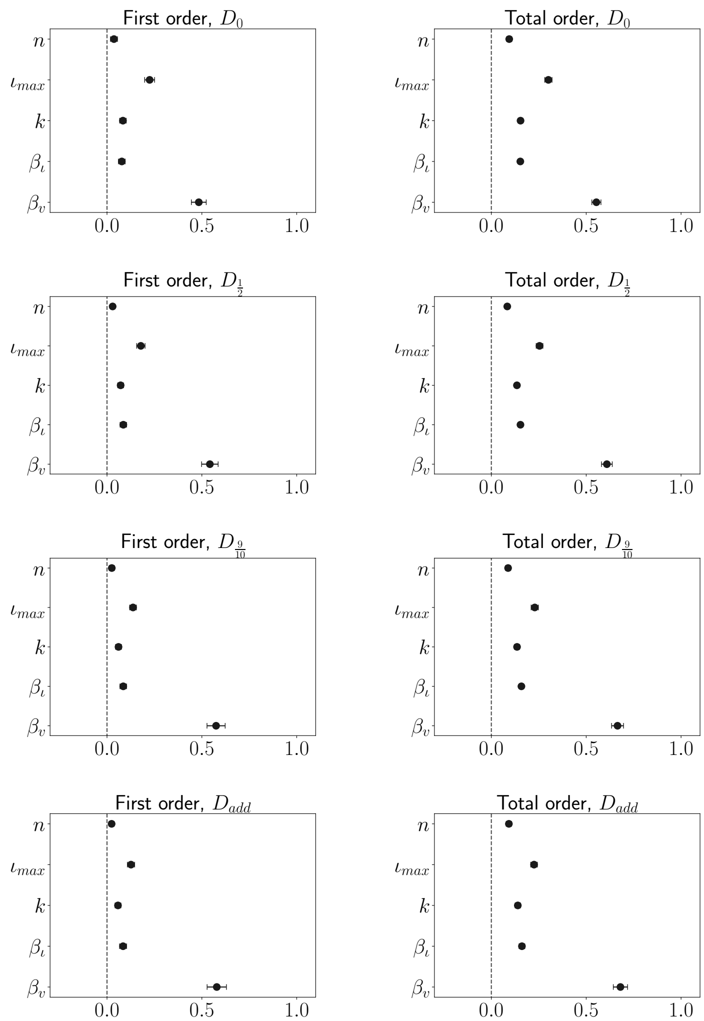

Figure 4.7: First and total order Sobol' indices for the effect of the model parameters n, $\iota_{\max}$, k, β_ι and β_v on the depth of the profiles generated by the model, where depth is measured as 0-viscous (i.e., the fraction of delegating voters), $\frac{1}{2}$-viscous, $\frac{9}{10}$-viscous and additive delegation depth. The analysis was run with $N = 4\,096$ for $d = 5$ model parameters, i.e., the model was run on $N \cdot (d + 2) = 28\,672$ samples from the input space. Error bars represent 95%-confidence intervals.

4.3.4 Epistemic Performance and Delegation Depth

Having studied the effect of different model parameters on the epistemic performance and delegation depth of the profiles generated, we turn to the correlation between delegation depth and epistemic performance. That is, we test whether the delegation depth of a profile can be used as a predictor of epistemic performance, in case we only have access to the liquid profile and not to the true preferences of all voters.

To perform this analysis, we generated a sample of the input space of our model. In order to guarantee that this sample is as representative as possible of the whole input space, we used a Saltelli sample (with $N = 128$ and thus $N \cdot (d + 2) = 896$ different combinations of input values). These points in the input space can be seen as different contexts in which voters generate their delegation structures. For each point in the input space, we generated a profile and computed our different measures of epistemic performance and delegation depth. The correlation between these values was computed using the correlation coefficient of Pearson [1895] and can be seen in Figure 4.8.

We note again that the plots for different measures of delegation depth look very similar. The correlation between delegation depth and expressive loss is quite strong with a correlation coefficient of about 0.6 (out of 1). This matches our hypothesis from Section 4.1: longer delegation chains in a liquid profile (or more precisely, deeper delegation structures according to our measures of delegation depth) tend to lead to less accurate representations of voters' true preferences in the proxy profile.

We further see that the correlation between delegation depth and systematic loss is moderate to weak at a correlation coefficient of 0.3. The weak correlation can be explained by the intuition that large expressive losses tend to lead to large systematic losses, but the inaccuracies of voters' ballots average out against each other as enough voters express their opinion inaccurately in varying directions.

Finally, the correlation between delegation depth and social welfare is very weak at a correlation coefficient of 0.1. An explanation could be that the collective decision of uniformly distributed voters is always close to $\frac{1}{2}$, and uniformly distributed voters are always at roughly the same average distance from $\frac{1}{2}$. Indeed, we see that the social welfare is around $\frac{1}{4}$ for all values of the delegation depth, which is exactly the expected distance between $\frac{1}{2}$ and a uniformly random position on $[0, 1]$.

To verify that the correlation coefficient between epistemic performance and delegation depth does not strongly depend on the measure of delegation depth that we choose, we computed the correlation coefficients for the additive delegation depth, and the α-viscous delegation depth for 20 different values of α. For each measure, we repeated the computation 10 times and calculated the median correlation coefficient. In Figure 4.9, we see that the median correlation coefficient for each measure of delegation depth is almost equal and the statistical errors are very small. We conclude that indeed, it does not make an important difference which measure of delegation depth we choose.

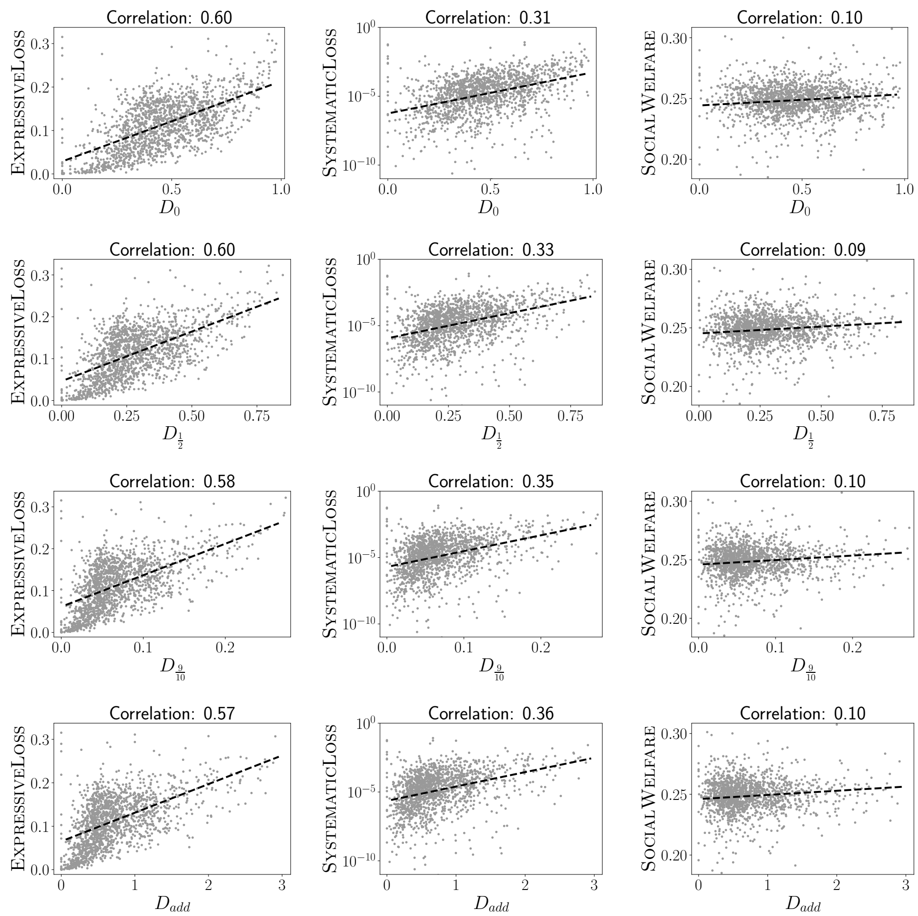

Figure 4.8: Correlation between the delegation depth of the profiles generated by our model and their epistemic performance. Each data point represents a single profile generated in a random context, i.e., a random combination of model parameters. The input space was sampled using a Saltelli sequence with $N = 128$ for $d = 5$ model parameters, i.e., $N \cdot (d + 2) = 896$ profiles were generated per plot. The correlation coefficients given are Pearson correlation coefficients [Pearson, 1895]. Dotted lines are linear fits to the data. To enhance visibility, up to 10% of outlying data points might fall outside the margins of each plot, and systematic loss was plotted on a logarithmic scale.

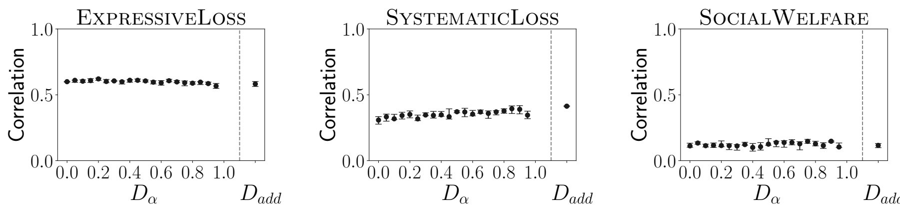

Figure 4.9: Analysis of the correlation between epistemic performance and different measures of delegation depth (D_α for $\alpha \in [0, 1)$ and D_{add}). Each data point is the median over 10 correlation coefficients, each computed on a Saltelli sample with $N = 128$. Error bars represent the first and third quartile.

This is a surprising result, since the 0-viscous delegation depth does not take into account the length of delegation chains, but simply counts the number of delegating voters; and the 1-viscous delegation depth is constantly equal to 0, which implies that the α-viscous delegation depth is arbitrarily close to zero for almost all delegation structures as α goes to 1. That is, the α-viscous delegation depth is a poor measure of delegation depth for very small or very large α, and we would expect these poor measures not to correlate strongly with epistemic performance. However, the geometric distribution might offer a possible explanation again. If the formation of delegation structures can be approximated as a geometric process, then the probability that a voter delegates (i.e., the 0-viscous delegation depth) contains all information needed to model the process. Thus, 0-viscous delegation depth might correlate with epistemic performance as strongly as any other measure of delegation depth, as long as the latter measure of delegation depth is not constant.

4.4 Discussion

We can draw a number of interesting conclusions from our model. Firstly, all measures of delegation depth that we defined (with the exception of 1-viscous delegation depth, which is constantly zero and was thus not analyzed) behave very similarly on the models we generated. They react almost identically to varying model parameters and correlate equally strongly with all measures of epistemic performance we defined. This shows a certain robustness of the model, but also makes it impossible to select a 'best' measure of delegation depth when using it to break ties in liquid aggregation.

Initially, it might be surprising that 0-viscous delegation depth predicts epistemic performance as accurately as the other measures of delegation depth, since 0-viscous delegation depth only takes into account the number of direct voters and not the actual delegation structure; it seems to contain less information than the other measures of delegation depth. However, if we approximate the process of profile formation as a geometrically distributed process, we see that the 0-viscous delegation depth contains exactly the right information to characterize the whole process: the probability that a voter decides to delegate. Therefore, it correlates equally strongly with epistemic performance as any other reasonable measure of delegation depth.

Secondly, delegation depth correlates significantly with epistemic loss. This could be used as an argument to place more trust in delegation structures which are less deep, and thus break ties between supporters and rejecters of a policy proposal in favor of the party with the most 'shallow' delegation structure. Importantly, although the correlation between systematic loss and delegation depth is weak, it is clearly positive and thus shallow structures do not tend to have a negative effect in terms of systematic loss. However, shallow delegation structures do exhibit slightly lower social welfare, but this correlation is almost negligible.

Finally, the number of voters does not influence any of our output measures, except for systematic loss (which can be understood as a consequence of Condorcet's jury theorem). This allows us to use our measures to analyze profiles of any size, at least up to 500 voters (which is the maximal number of voters we considered).

Of course, it is important to note that we base our conclusions on a strongly simplified model of voter behavior. Empirical evidence should be gathered to confirm the results of our work. Furthermore, epistemic loss is only a single argument to favor shallow delegation structures over deep delegation structures. Thus, we do not recommend using, e.g., viscous democracy as proposed by Boldi et al. [2011] (in which not only tie-breaking, but also computing actual voting weight is done using a dampening factor α per delegation step) on the basis of this single argument (although other appealing arguments for viscous democracy could of course be formulated, and have been formulated by its authors).

Apart from studying profile formation, our model can also be used to generate profiles for other purposes. For example, in Section 5.3 we will use profiles generated by this model to numerically evaluate the performance of liquid aggregation rules. Furthermore, in combination with empirical data, one can attempt to discover biases in voter behavior, by matching real-world liquid profiles with liquid profiles generated by a model with different values for biases β_ι and β_v, and other biases.

Chapter 5

A Structural Ranked Agenda Rule

In Chapters 2 and 3, we defined a judgment aggregation formalism for the liquid democracy setting, axiomatically studied some impossibilities in designing its aggregation rules, and identified the ranked agenda rule as a reasonable, though irresolute, method to arrive at complete and consistent collective decisions. In Chapter 4, we computationally studied the relationship between delegation structure and epistemic performance of a liquid democracy profile, and concluded that deep delegation structures signify relatively poor epistemic accuracy. In this chapter, we combine these two perspectives and propose a refinement of the ranked agenda rule, which takes into account the delegation structure of a profile beyond the common proxy profile approach.

In Section 5.1, we define our structural ranked agenda rule, along with the more radical 'viscous democracy' approach. In order to compare our aggregation rules to each other, and to Kemeny's rule as a benchmark, we extend our computational model of profile formation to aggregation problems with multiple issues and an integrity constraint in Section 5.2. In Section 5.3, we present the results of the numerical evaluation of our aggregation rules.

5.1 Definition and Normative Properties

In Section 4.3, we learned that all of our measures of delegation depth correlate comparably strongly with the epistemic performance of the underlying profile. Therefore, we focus our attention in this chapter on the simplest measure: 0-viscous delegation depth, i.e., the fraction of delegating voters. This simple parameter seems to dictate any measure of delegation depth as well as the epistemic performance of a profile (possibly because the computational model resembles a geometrically distributed process of delegation) and can simultaneously be meaningfully interpreted as the fraction of voters who are (or feel) capable of determining their own optimal direct vote.

For normative reasons, our main proposal does not reduce the weight of voters who delegate to a long chain of peers: this would violate the principle of 'one man, one vote'. Instead, we conservatively propose to break ties in the ranked agenda rule according to delegation depth. Given a liquid aggregation profile $\boldsymbol{B} \in \mathcal{B}^\Phi$ (in the sense of Definition 2.1.3, not the sense of Section 4.2.1) and a literal $\ell \in \Phi$, we define their *(0-viscous) delegation depth* as

$$D_0(\boldsymbol{B}, \ell) = \frac{|\{i \in \mathcal{N}_{\boldsymbol{B}} \mid B_i(\ell) \in \mathcal{N}_{\boldsymbol{B}}\}|}{|\mathcal{N}_{\boldsymbol{B}}|}.$$

Note that this is equivalent to the definition of α-viscous delegation depth in Section 4.2.3 for $\alpha = 0$. We can now define our structural ranked agenda rule as follows.

Definition 5.1.1 (Structural Ranked Agenda Rule)**.** For any profile $\boldsymbol{B} \in \mathcal{B}^\Phi$, let $\succeq_{\boldsymbol{B}}^{D_0}$ be the binary relation on Φ where for $\ell_1, \ell_2 \in \Phi$, we have $\ell_1 \succeq_{\boldsymbol{B}}^{D_0} \ell_2$ if and only if

$$n_{\boldsymbol{B}}(\ell_1) > n_{\boldsymbol{B}}(\ell_2), \text{ or } n_{\boldsymbol{B}}(\ell_1) = n_{\boldsymbol{B}}(\ell_2) \text{ and } D_0(\boldsymbol{B}, \ell_1) \geq D_0(\boldsymbol{B}, \ell_2).$$

Let $\mathcal{LO}(\succeq_{\boldsymbol{B}}^{D_0})$ be the set of linear orders on Φ that are compatible with $\succeq_{\boldsymbol{B}}^{D_0}$. The *structural ranked agenda rule* is the liquid aggregation rule $\textsc{ra}_{\text{struc}}$ generated by the following process.

- Given profile $\boldsymbol{B}$ and integrity constraint Γ, initialize $\textsc{ra}_{\text{struc}}(\boldsymbol{B}, \Gamma) = \emptyset$.
- For each linear order $\succeq \in \mathcal{LO}(\succeq_{\boldsymbol{B}}^{D_0})$, do the following.
  - Initialize $J = \emptyset$.
  - Iteratively, in the order $\succeq$, consider a literal $\ell \in \Phi$. If $J \cup \{\Gamma\} \models {\sim}\ell$, add ${\sim}\ell$ to J. Otherwise, add ℓ to J.
  - After considering all $\ell \in \Phi$, add J to $\textsc{ra}_{\text{struc}}(\boldsymbol{B}, \Gamma)$.

Since the relation $\succeq_{\boldsymbol{B}}^{D_0}$ is compatible with the relation $\succeq_{\boldsymbol{B}}$ of the original ranked agenda rule (see Definition 3.1.1), the structural ranked agenda rule is indeed a refinement of the ranked agenda rule which breaks ties according to delegation depth. Since 0-viscous delegation depth takes values in the interval $[0, 1) \subseteq \mathbb{R}$, we can alternatively define the *liquid support* for a literal $\ell \in \Phi$ in profile $\boldsymbol{B} \in \mathcal{B}^\Phi$ as the sum $s_{\boldsymbol{B}}(\ell) = n_{\boldsymbol{B}}(\ell) + D_0(\boldsymbol{B}, \ell)$, and view $\succeq_{\boldsymbol{B}}^{D_0}$ as a (weak) ordering of the agenda Φ by liquid support of the literals.

In light of Section 3.2, we also give a functional definition of the structural ranked agenda rule.

Definition 5.1.2 (Functional Structural Ranked Agenda Rule)**.** For any profile $\boldsymbol{B} \in \mathcal{B}^\Phi$ and integrity constraint $\Gamma \in \mathcal{L}_{\text{sat}}$, let $\succeq^{\boldsymbol{B}}_{\text{struc}}$ be the binary relation on $\mathcal{J}(\Gamma)$ defined by

$$J \succeq^{\boldsymbol{B}}_{\text{struc}} J' \text{ if and only if } \min_{\ell \in J \setminus J'} s_{\boldsymbol{B}}(\ell) \geq \min_{\ell \in J' \setminus J} s_{\boldsymbol{B}}(\ell)$$

for $J, J' \in \mathcal{J}(\Gamma)$, where the minimum over the empty set is defined to equal ∞. The *functional structural ranked agenda rule* is the liquid aggregation rule $\text{RA}^f_{\text{struc}}$ which maps any profile $\boldsymbol{B} \in \mathcal{B}^\Phi$ and integrity constraint $\Gamma \in \mathcal{L}_{\text{sat}}$ to

$$\text{RA}^f_{\text{struc}}(\boldsymbol{B}, \Gamma) = \{J \in \mathcal{J}(\Gamma) \mid \forall J' \in \mathcal{J}(\Gamma)\colon J \succeq^{\boldsymbol{B}}_{\text{struc}} J'\}.$$

Since liquid support $s_{\boldsymbol{B}}(\ell)$ gives rise to a weak order on Φ (just like majority margins $n_{\boldsymbol{B}}(\ell)$ do), the proof of Proposition 3.2.4 translates directly to an analogous proposition for the structural ranked agenda rule, showing the equivalence of Definitions 5.1.1 and 5.1.2.

Proposition 5.1.3. *For any profile $\boldsymbol{B} \in \mathcal{B}^\Phi$ and integrity constraint $\Gamma \in \mathcal{L}_{sat}$, we have*

$$\textit{RA}_{\text{struc}}(\boldsymbol{B}, \Gamma) = \textit{RA}^f_{\text{struc}}(\boldsymbol{B}, \Gamma).$$

Likewise, the arguments for Proposition 3.1.3 translate directly to the structural ranked agenda rule, proving that it has the following normative properties.

Proposition 5.1.4. *The structural ranked agenda rule on the universal domain is language-independent, collectively complete, complement-free and consistent, anonymous, contextually neutral, and contextually unbiased. It is not resolute, static, neutral, unbiased, or independent.*

Since the axioms which were defined in Section 2.2.3 and which characterize the ranked agenda rule, are formulated in terms of majority margins (and not liquid support), the structural ranked agenda rule does not satisfy all of these properties (but it does satisfy weak Condorcet consistency, weak homogeneity and monotonic consistency by similar arguments as for the ranked agenda rule). If we defined analogous axioms in terms of liquid support, the structural ranked agenda rule would satisfy these axioms. But such axioms would strongly depend on our specific method of breaking ties, and are thus arguably not suited as general normative principles.

As stated before, the structural ranked agenda rule is a conservative proposal which only breaks ties in the ranked agenda rule. A more radical approach by Boldi et al. [2011], which we briefly discussed in Section 4.2.3, reduces the voting weight of voters by a factor $\alpha \in [0, 1]$ for each step in their delegation path. For $\alpha = 0$, this approach amounts to ignoring all voters who do not vote directly, and for $\alpha = 1$, this approach is equivalent to the original ranked agenda rule. Thus, by decreasing the value of α, we make the aggregation rule more 'viscous' (i.e., less 'liquid'), which is why the approach is referred to as 'viscous democracy'. In order to compare our structural ranked agenda rule to viscous democracy, we define an aggregation rule which computes the viscous voting weights of all voters, and runs the ranked agenda algorithm prioritizing the literals by total viscous voting weight of the supporting voters.

Consider any viscosity factor $\alpha \in [0,1]$, any profile $\boldsymbol{B} \in \mathcal{B}^\Phi$ and any integrity constraint $\Gamma \in \mathcal{L}_{\text{sat}}$. We define the α-viscous voting weight for each voter $i \in \mathcal{N}_{\boldsymbol{B}}$ and literal $\ell \in \Phi$ as

$$w_i^\alpha(\ell) = 1 + \sum_{\{j \in \mathcal{N}_{\boldsymbol{B}} \mid B_j(\ell) = i\}} \alpha \cdot w_j^\alpha(\ell),$$

where the empty sum is defined to equal 0. Let

$$v_{\boldsymbol{B}}^\alpha(\ell) = \sum_{\{i \in \mathcal{N}_{\boldsymbol{B}} \mid B_{\text{rep}_{\boldsymbol{B},\ell}(i)}(\ell) = +\}} w_i^\alpha$$

be the *viscous support* of literal $\ell \in \Phi$ under profile $\boldsymbol{B} \in \mathcal{B}^\Phi$. Then the *$\alpha$-viscous ranked agenda rule* is defined as follows.

Definition 5.1.5 (Viscous Ranked Agenda Rule)**.** Let $\alpha \in [0,1]$. For any profile $\boldsymbol{B} \in \mathcal{B}^\Phi$, let $\succeq_{\boldsymbol{B}}^\alpha$ be the binary relation on Φ where for $\ell_1, \ell_2 \in \Phi$, we have $\ell_1 \succeq_{\boldsymbol{B}}^\alpha \ell_2$ if and only if $v_{\boldsymbol{B}}^\alpha(\ell_1) \geq v_{\boldsymbol{B}}^\alpha(\ell_2)$. Let $\mathcal{LO}(\succeq_{\boldsymbol{B}}^\alpha)$ be the set of linear orders on Φ that are compatible with $\succeq_{\boldsymbol{B}}^\alpha$. The *$\alpha$-viscous ranked agenda rule* is the liquid aggregation rule RA_α generated by the following process.

- Given profile $\boldsymbol{B}$ and integrity constraint Γ, initialize $\text{RA}_\alpha(\boldsymbol{B}, \Gamma) = \emptyset$.
- For each linear order $\succeq \in \mathcal{LO}(\succeq_{\boldsymbol{B}}^\alpha)$, do the following.
  - Initialize $J = \emptyset$.
  - Iteratively, in the order $\succeq$, consider a literal $\ell \in \Phi$. If $J \cup \{\Gamma\} \models {\sim}\ell$, add ${\sim}\ell$ to J. Otherwise, add ℓ to J.
  - After considering all $\ell \in \Phi$, add J to $\text{RA}_\alpha(\boldsymbol{B}, \Gamma)$.

Clearly, we could define a functional version of the viscous ranked agenda rule again, and show that it is equivalent to its algorithmic definition. Note that the viscous ranked agenda rule satisfies the same properties from Proposition 3.1.3 as the ranked agenda rule (by analogous arguments), and from the axioms of Section 2.2.3, it only satisfies weak homogeneity and monotonic consistency.[1]

5.2 Profile Formation for Multiple Issues

In Section 4.2, we defined a model of profile formation for liquid democracy. However, our model only generates profiles for single-issue agendas. In this section, we extend our model to generate rational liquid aggregation profiles

[1]Unlike the structural ranked agenda rule, the α-viscous ranked agenda rule for $\alpha \neq 1$ does not satisfy weak Condorcet consistency, since the total viscous weight in favor of a proposition can be larger than the total viscous weight against it, even if the proposition receives strict majority support in terms of majority margins.

on multiple issues, which are logically connected through a satisfiable integrity constraint.

Consider an agenda $\Phi = \{p_1, \neg p_1, \ldots, p_m, \neg p_m\}$ containing $m \in \mathbb{N}$ propositional letters and their negations. A liquid aggregation profile $\boldsymbol{B}$ for this agenda essentially consists of m separate single-issue profiles on some common set of voters $\mathcal{N}$. However, whereas a liquid aggregation profile in the sense of Definition 2.1.3 either accepts a propositional letter or rejects it, recall that our model of single-issue profile formation allows direct votes to take any real value between 0 and 1. Thus, a multi-issue ballot by voter $i \in \mathcal{N}$ in our computational model corresponds to a point $\boldsymbol{b}_i'$ in the space $[0,1]^k \subseteq \mathbb{R}^k$ (where $k \in \{0, \ldots, m\}$ denotes the number of issues on which voter i votes directly), along with $m-k$ delegations. To convert a computationally generated multi-issue profile to a proper liquid aggregation profile (see Example 5.2.1), we place all complete and complement-free judgment sets J for agenda Φ in the space $[0,1]^m$, where the ith coordinate of J is 1 if and only if $p_i \in J$, and the ith coordinate of J is 0 if and only if $\neg p_i \in J$. Consequently, for each voter $i \in \mathcal{N}$, we find the complete and complement-free judgment set $J^* \subseteq \Phi$, which minimizes the k-dimensional Euclidean distance between $\boldsymbol{b}_i'$ and J^*, where the position of J^* in $[0,1]^k$ is simply the projection the position of J^* in $[0,1]^m$ along the issues on which voter i votes directly. The k direct votes of voter i in liquid aggregation profile $\boldsymbol{B}$ are then defined to follow the judgments of J^*, and the $m-k$ delegations are identical to the delegations in the corresponding $m-k$ single-issue profiles.

If we further introduce an integrity constraint $\Gamma \in \mathcal{L}_{\text{sat}}$ and require the profile to be rational (as proposed in Section 1.4 and formalized in Section 2.2.1), we computationally generate the single-issue profiles identically, but when converting to a multi-issue liquid aggregation profile, we only consider the complete and complement-free judgment sets J which are consistent with the integrity constraint.

The following example illustrates the conversion of computationally generated single-issue profiles to a multi-issue liquid aggregation profile, from the perspective of some voter $i \in \mathcal{N}$.

Example 5.2.1. Consider some set of voters $\mathcal{N}$ (containing at least two voters) and the agenda $\Phi = \{p_1, \neg p_1, p_2, \neg p_2, p_3, \neg p_3\}$ with $m = 3$. In order to generate a liquid aggregation profile $\boldsymbol{B}$ for agenda Φ, we use the model of Section 4.2 to generate three single-issue profiles for voters $\mathcal{N}$. Suppose voter $i \in \mathcal{N}$ has m-dimensional position $\pi_i = (0.5, 0.6, 0.1)$ and ballot B_i' with $B_i'(p_1) = 0.6$, $B_i'(p_2) = 0.3$ and $B_i'(p_3) = j$ for some voter $j \in \mathcal{N} \setminus \{i\}$ (see the gray elements in Figure 5.1). In other words, voter i votes directly on p_1 and p_2, and delegates her vote on p_3.

To convert the real-valued ballot of voter i to a proper liquid aggregation ballot, we consider all complete and complement-free judgment sets J for Φ. Since voter i delegates her vote on p_3 to another voter, we are only interested in the judgments in J on p_1 and p_2. In Figure 5.1a, we see that the judgment sets J^* with $\{p_1, \neg p_2\} \subseteq J^*$ have minimal distance to the direct votes of voter i.

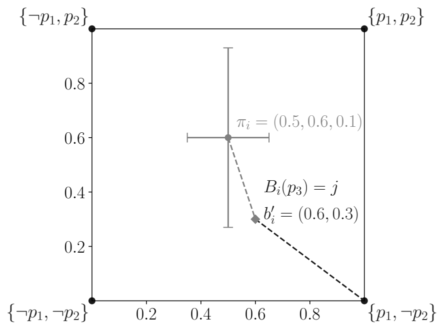

(a) No integrity constraint.

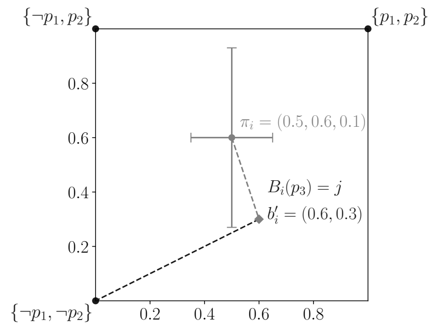

(b) Integrity constraint $p_1 \to p_2$.

Figure 5.1: Generation of a 3-issue liquid aggregation ballot, where voter i delegates her vote on p_3. The horizontal axes represent her judgment on p_1 and the vertical axes represent her judgment on p_2.

Thus, we set $B_i(p_1) = +$ and $B_i(p_2) = -$. Furthermore, we leave the delegation of voter i intact and set $B_i(p_3) = j$.

Note that in this conversion, we did not consider any integrity constraint. If we introduce integrity constraint $\Gamma = p_1 \to p_2$ and require that the liquid aggregation profile $\boldsymbol{B}$ is rational (i.e., the direct votes of all voters are consistent with the integrity constraint), then the judgment sets J with $\{p_1, \neg p_2\} \subseteq J$ are not allowed. As a consequence, the judgment sets $J^* \in \mathcal{J}(\Gamma)$ with $\{\neg p_1, \neg p_2\} \subseteq J^*$ have minimal distance to the direct votes of voter i (see Figure 5.1b). Therefore, in case we require rationality, we set $B_i(p_1) = -$, $B_i(p_2) = -$ and $B_i(p_3) = j$. △

5.3 Numerical Evaluation of Four Rules

Using our model for liquid aggregation profile formation, we numerically compare the performance of four liquid aggregation rules: Kemeny's rule, the original ranked agenda rule (Definition 3.1.1), the structural ranked agenda rule (Definition 5.1.1) and the $\frac{1}{2}$-viscous ranked agenda rule (Definition 5.1.5). To do so, we randomly generated 1 000 propositional formulas in conjunctive normal form (containing at most 15 different propositional letters, at most 10 clauses and at most 5 literals per clause, and having at least two different satisfying truth assignments). Consequently, for each formula, we generated a Saltelli sequence of size $N = 3$ of the 5-dimensional parameter space of our single-issue profile formation model (with parameters $n \in \{2, \ldots, 500\}$, $\iota_{\max} \in [0, 1]$, $k \in \{1, \ldots, 100\}$, $\beta_\iota \in [0, 2]$ and $\beta_v \in [0, 2]$; see Sections 4.2 and 4.3.1) and used it to generate $N \cdot (d + 2) = 56$ rational profiles with the formula as the integrity constraint.

Given a collective decision $J \in \mathcal{J}(\Gamma)$, we define the satisfaction of voter $i \in \mathcal{N}$ to be inversely proportional to the Euclidean distance between J and the position π_i of voter i. In other words, the satisfaction s_i of voter i is

$$s_i = \frac{1}{d(J, \pi_i)} = \frac{1}{\sqrt{\sum_{p_j \in J}(1 - (\pi_i)_j)^2 + \sum_{\neg p_j \in J}((\pi_i)_j)^2}}.$$

For each of the 56 000 profiles generated, we used a brute force algorithm to find the collective decision(s) that maximized average voter satisfaction. Furthermore, we computed the collective decisions generated by our four aggregation rules.

In Figure 5.2, we see for each pair $\mathcal{F}, \mathcal{F}'$ of aggregation rules how often $\mathcal{F}(\boldsymbol{B})$ was a subset of $\mathcal{F}'(\boldsymbol{B})$, where $\boldsymbol{B}$ ranges over all 56 000 profiles generated. Since the structural ranked agenda rule is a refinement of the ranked agenda rule, we see that the collective decision under the structural ranked agenda rule is indeed a subset of the collective decision under the ranked agenda rule in 100% of cases. Furthermore, it seems that Kemeny's rule, the ranked agenda rule and the structural ranked agenda rule are more similar to each other than to the optimal rule or the viscous ranked agenda rule. This is an interesting observation, since the decision mechanisms behind the structural ranked agenda rule and the viscous ranked agenda rule are intuitively similar, but apparently produce rather different outcomes. Finally, we note that no aggregation rule selects the optimal collective decision in more than 11% of cases. Note that the latter observation might be an artifact of our model of profile formation and our definition of voter satisfaction: possibly, voters often submit ballots that are rather different from their optimal ballot. For example, in Figure 5.1 we see that the ignorance of the voter is large, causing her to vote for judgment

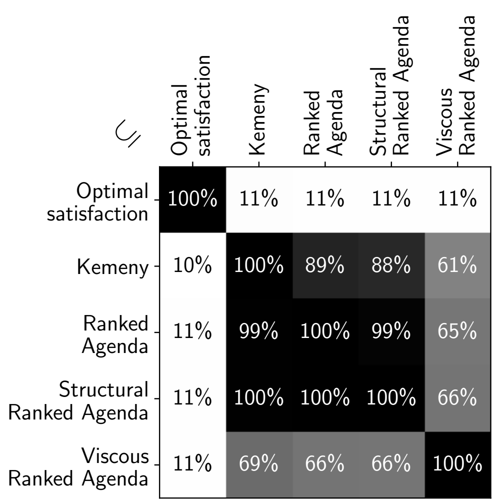

Figure 5.2: Frequency with which each aggregation rule returns a subset of another aggregation rule.

sets $\{p_1, \neg p_2\}$ or $\{\neg p_1, \neg p_2\}$, whereas judgment sets $\{\neg p_1, p_2\}$ and $\{p_1, p_2\}$ are closest to her actual position. Although the judgment set that is closest to the position of a voter is *most* likely to be selected, deviations may occur frequently. And if the profile approximates the voters' preferences poorly, an aggregation rule cannot consistently arrive at the optimal collective decision.

Figure 5.3 plots three normative measures of the aggregation rules. Firstly, we see the error with which each rule approximates the optimal average voter satisfaction. The error is defined as the percentage (averaging over all profiles) by which the voter satisfaction under a given aggregation rule is smaller than the optimal voter satisfaction. Secondly, we see the frequency with which each rule was irresolute (i.e., returned more than one possible collective decision). And thirdly, we see how many possible decisions each rule returned on average, in case it returned more than one profile. Note that the latter measure is always greater than (or equal to) 2, since we only take into account the profiles for which the rule is irresolute.

In the figure, we see that all aggregation rules approximate the optimal voter satisfaction with an error of around 0.85%. We consider the differences between the approximation errors negligible, even if they are statistically significant. On the other hand, we clearly see a difference in irresoluteness between the rules: Kemeny's rule returns multiple possible collective decisions in almost 12% of cases, whereas the different ranked agenda rules rarely return multiple possible collective decisions. Moreover, the structural ranked agenda rule is slightly more resolute than the viscous ranked agenda rule, which is in turn more resolute than the ranked agenda rule. Finally, we see that if one of the ranked agenda rules is irresolute, it almost always returns two possible collective decisions, whereas Kemeny's rule returns an average of 2.5 possible collective decisions, in case it is irresolute.

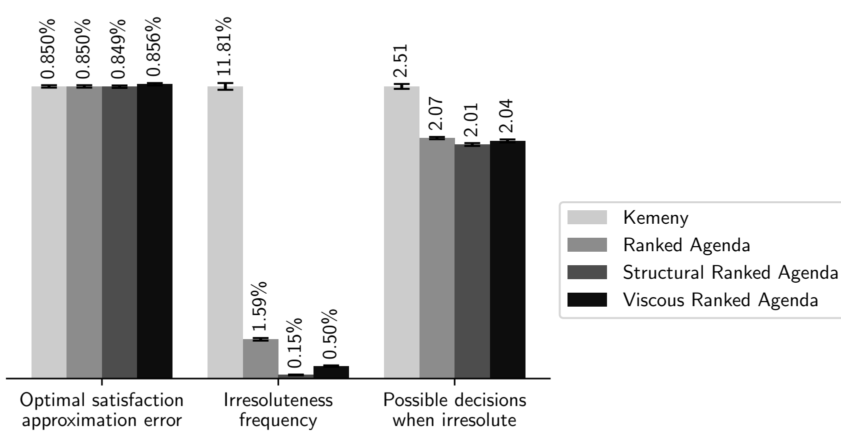

Figure 5.3: Three normative measures of Kemeny's rule and three variants of the ranked agenda rule. Error bars denote the standard error of the mean.

For a more detailed understanding of how well each aggregation rule approximates the optimal voter satisfaction, we performed an OFAT analysis (see Section 4.3.1) of the optimal satisfaction approximation error. However, in Figure 5.4 we see that none of the model parameters significantly influence the approximation error (although the number of voters n and the maximal ignorance $\iota_{\max}$ seem to influence the approximation error weakly, but consistently). Possibly, this is an adverse effect of the method in which we generate profiles and evaluate voter satisfaction. Since voters' positions are uniformly distributed along the axis of every issue, the average position is very close to the midpoint of the opinion space. Therefore, all judgment sets (i.e., corners of the opinion space) are roughly equally satisfactory to the average voter. By reducing the number of voters, we allow the average position of the voters to have a larger variance, thus reducing this issue.

In Figure 5.5, we reduced the default value of n from 250 to 10, and the parameter range from $\{1, \ldots, 500\}$ to $\{1, \ldots, 20\}$. Firstly, we note that the approximation error increases as the maximal ignorance $\iota_{\max}$ increases, and as the number of issues m increases. Both are easily explained: if voters are less ignorant, their direct votes are closer to their actual positions, thus increasing the accuracy of the aggregation rule; and if there are fewer issues to vote on, there are fewer complete and consistent judgment sets, thus increasing the frequency with which the aggregation rules select the optimal collective decision. Furthermore, if we compare Figure 5.4 to Figure 5.5, we see that the approximation error is much larger in the latter figure (between 0% and 10%, instead of between 0% and 2%), and especially so for very small numbers of voters ($n < 5$). This confirms our conjecture in the previous paragraph, that for large numbers of uniformly distributed voters, all judgment sets are roughly equally distant from the average voter, allowing all four aggregation rules to approximate the optimal collective judgment very accurately. Thus, in order to compare the performance of our aggregation rules more thoroughly, we need to generate our profiles over a less uniformly distributed set of voters, or gather empirical voting data.

In line with this reasoning, it is important to note that all results in this section depend strongly on the integrity constraints generated by our randomized algorithm. Since real-world applications of liquid democracy might consider very specific types of integrity constraints, the aggregation rules might perform very differently in practice, or between different applications.

From our experiments, we conclude that none of the rules we analyzed return the optimal collective judgment very often. However, all rules approximate the optimal voter satisfaction rather accurately with an error of around 0.85%. Moreover, the structural ranked agenda rule and the viscous ranked agenda rule are considerably more resolute than the ranked agenda rule, and much more resolute than Kemeny's rule. A more detailed analysis of the approximation error of our aggregation rules can be obtained by generating less uniformly distributed voters, or collecting empirical voting data.

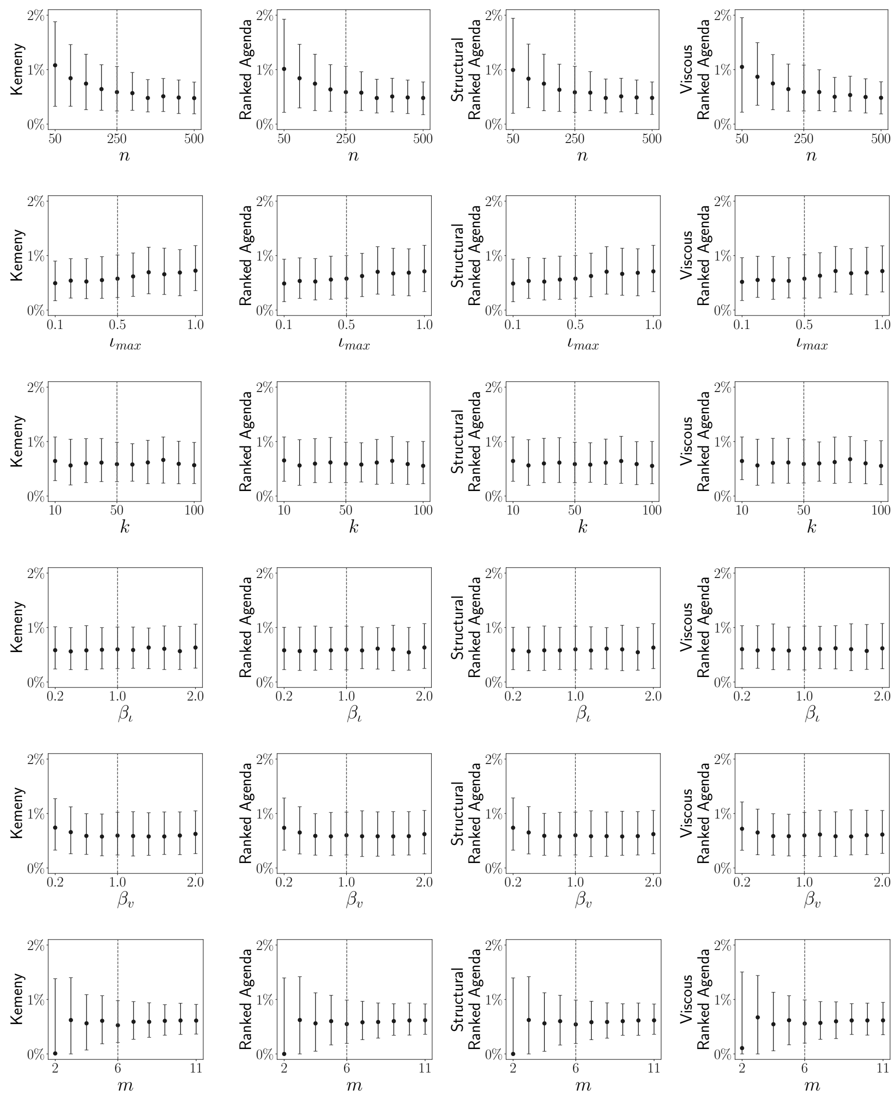

Figure 5.4: OFAT analysis of the effect of the model parameters n, $\iota_{\max}$, k, β_ι, β_v and m on the optimal satisfaction approximation error of four liquid aggregation rules. Each data point is the median output over 1 000 profiles with each a unique randomly generated integrity constraint over m propositional letters. Error bars represent the first and third quartile. Dotted lines indicate the default value of each parameter.

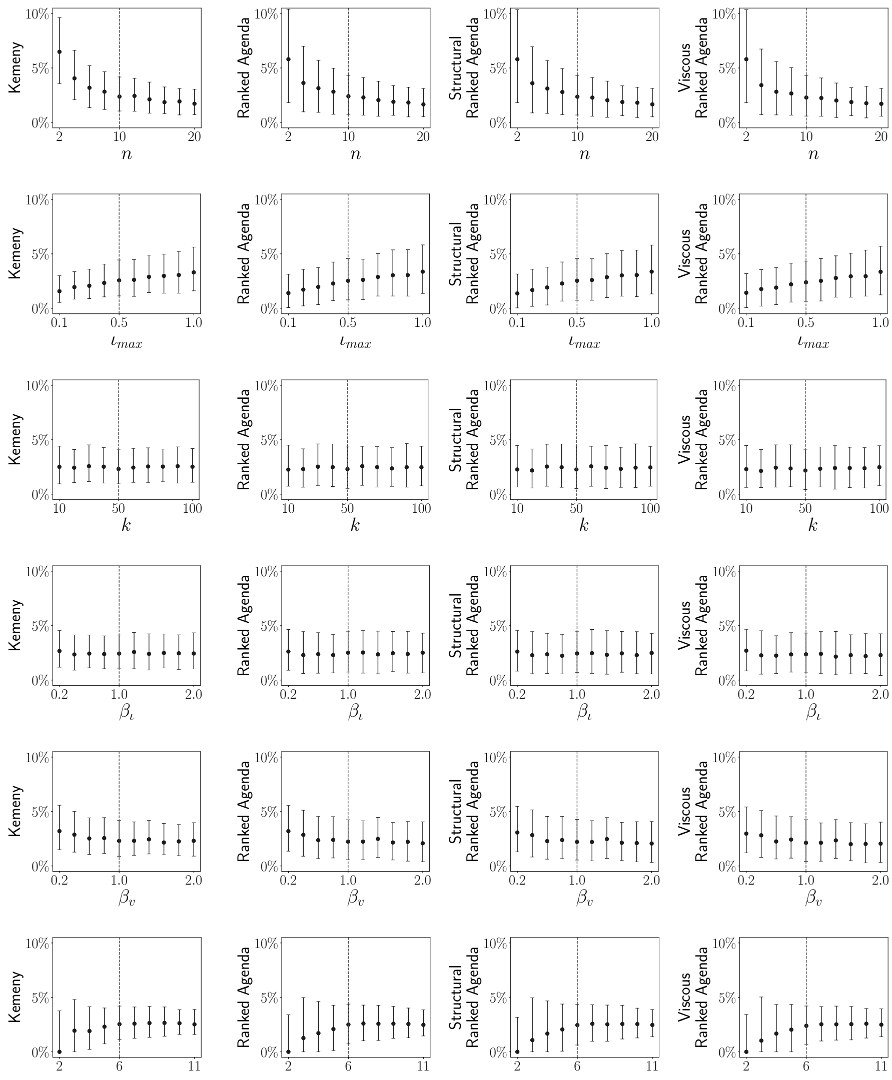

Figure 5.5: OFAT analysis of the optimal satisfaction approximation error of four liquid aggregation rules for small numbers of voters n. Each data point is the median output over 1 000 profiles. Error bars represent the first and third quartile. Dotted lines indicate the default value of each parameter.

Chapter 6

Conclusion

The main objective of this thesis was to design a normatively appealing liquid aggregation rule that guarantees complete and consistent collective judgments. Our approach to this problem consisted of three distinguishable components: the generalization of elements from formal preference and judgment aggregation theory to the liquid democracy setting, the computational study of epistemic performance of aggregation rules in connection to delegation structure, and the design and computational evaluation of a structural version of the ranked agenda rule for liquid democracy.

In the first component (Chapters 2 and 3), we defined a formal model of liquid judgment aggregation, studied some of its impossibility theorems and identified the ranked agenda rule as a computationally tractable and normatively appealing liquid aggregation rule. Our main contributions to the existing body of research in these chapters include the definition of irresolute and/or contextual versions of the most important normative axioms of judgment aggregation theory, which preserve the normative principle behind the formal axioms and (therefore) preserve the most important impossibility results. Furthermore, we showed that the axiomatic characterization by Lamboray [2009b] of the ranked pairs rule as the largest monotone and prudent preference aggregation rule translates directly to the liquid judgment aggregation setting.

In the second component (Chapter 4), we computationally studied the formation of liquid profiles and the resulting correlation between the depth of a profile's delegation structure and its epistemic performance. We concluded that deep delegation structures tend to approximate voters' preferences less accurately. Therefore, according to our randomly generated profiles, delegation depth can be used as a proxy for relatively unreliable liquid profiles. We hope that our computational model of voter behavior (and the general approach taken) can be used in the future to study the process of profile formation and to computationally compare different aggregation mechanisms.

In the final component (Chapter 5), we combined the formal study of the ranked agenda rule and the computational study of delegation structures to propose a structural version of the ranked agenda rule, which prioritizes issues based on the number of supporters, breaking ties by delegation depth. In a computational comparison between the structural ranked agenda rule and three other complete and consistent judgment aggregation rules (Kemeny's rule, the original ranked agenda rule and viscous democracy), all rules seemed to approximate the optimal average voter satisfaction equally well, while the structural ranked agenda rule was considerably more resolute than the other rules.

As argued in the introduction of this thesis, the concept of liquid democracy only has potential if it can avoid policy-inconsistencies. Aggregating judgments by means of a collectively consistent aggregation rule is one of the possible solutions to this problem. Although computational social choice theory teaches us that collectively complete and consistent aggregation rules are generally computationally intractable, irresolute, or both, we have seen that in practice, our structural ranked agenda rule can be computed efficiently (as long as the integrity constraint does not contain too many issues) and rarely returns more than one (and very rarely more than two) possible collective judgments. Therefore, it is an important rule to consider when designing real-world liquid aggregation instances.

We should note that our research leans heavily on a very simple model of voter behavior, and one in which voters' ideological positions are uniformly distributed. Furthermore, our integrity constraints were randomly generated. Therefore, our findings should be confirmed by a more realistic model of voter behavior in combination with integrity constraints obtained from real-world aggregation problems, or ideally by empirical voting data. Furthermore, although the structural ranked agenda rule can be applied to irrational and/or cyclic profiles, in the computational analysis we assumed our models to be rational and acyclic. Allowing for irrational voters, cyclic delegation and abstention possibly alters the results significantly.

Therefore, future work should focus on the empirical verification of our conclusions, ideally in the presence of irrational, cyclic or incomplete profiles. Furthermore, our work was inconclusive in deciding on a 'best' measure of delegation depth. Possibly, one of our measures of delegation depth (or some other measure on delegation structures) is a better predictor of epistemic performance or of the accurate approximation of the optimal collective decision than the fraction of casting voters, which we used in the structural ranked agenda rule. Finally, it would be interesting to formally study further (or other) refinements of the ranked agenda rule, such as the leximax rule. Whereas the ranked agenda rule is the largest *monotone* and prudent rule, other prudent rules might be characterized by similarly appealing axioms.

We conclude that the issue of policy-inconsistency in liquid democracy has a fair chance of being resolved by the clever design of aggregation mechanisms. The structural ranked agenda rule, for instance, is both a normatively appealing and an experimentally accurate method of arriving at satisfactory, consistent collective decisions. Having said that, this inquiry to liquid democracy has only left us "curiouser and curiouser" – in the words of Lewis Carroll [1865] – in search of the ideal voting method for the democratic ideal.

Bibliography

K. J. Arrow and H. Raynaud. *Social Choice and Multicriterion Decision-Making*. MIT Press, 1986.

J. Bartholdi, III, 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.

J. Behrens. The origins of liquid democracy. *The Liquid Democracy Journal*, 5: 7–17, 2017.

J. Behrens, A. Kistner, A. Nitsche, and B. Swierczek. *The Principles of LiquidFeedback*. Interaktive Demokratie e.V., 2014.

C. Blum and C. I. Zuber. Liquid democracy: Potentials, problems, and perspectives. *Journal of Political Philosophy*, 24(2):162–182, 2016.

G. Boella, L. Francis, E. Grassi, A. Kistner, A. Nitsche, A. Noskov, L. Sanasi, A. Savoca, C. Schifanella, and I. Tsampoulatidis. WeGovNow: A map based platform to engage the local civic society. In P.-A. Champin, F. Gandon, L. Médini, M. Lalmas, and P. G. Ipeirotis, editors, *Companion Proceedings of the 27th The Web Conference (WWW)*, pages 1215–1219. International World Wide Web Conferences Steering Committee, 2018.

P. Boldi, F. Bonchi, C. Castillo, and S. Vigna. Viscous democracy for social networks. *Communications of the ACM*, 54(6):129–137, 2011.

M. Brill. Interactive democracy. In E. André, S. Koenig, M. Dastani, and G. Sukthankar, editors, *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1183–1187. International Foundation for Autonomous Agents and Multiagent Systems, 2018.

M. Brill and F. Fischer. The price of neutrality for the ranked pairs method. In J. Hoffmann, B. Selman, and D. Fox, editors, *Proceedings of the 26th AAAI Conference on Artificial Intelligence*, pages 1299–1305. Association for the Advancement of Artificial Intelligence, 2012.

M. Brill and N. Talmon. Pairwise liquid democracy. In J. Lang, editor, *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 137–143. Association for the Advancement of Artificial Intelligence, 2018.

M. Brill, T. Delemazure, A.-M. George, M. Lackner, and U. Schmidt-Kraepelin. Liquid democracy with ranked delegations. In K. Sycara, V. Honavar, and M. T. J. Spaan, editors, *Proceedings of the 36th AAAI Conference on Artificial Intelligence*, pages 4884–4891. Association for the Advancement of Artificial Intelligence, 2022.

L. Carroll. *Alice's Adventures in Wonderland*. Macmillan Publishers, 1865.

Z. Christoff and D. Grossi. Binary voting with delegable proxy: An analysis of liquid democracy. In J. Lang, editor, *Proceedings of the 16th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, pages 134–150. Electronic Proceedings in Theoretical Computer Science, 2017.

R. Colley, U. Grandi, and A. Novaro. Unravelling multi-agent ranked delegations. *Autonomous Agents and Multi-Agent Systems*, 36(1), 2022.

N. Condorcet. *Essai sur l'Application de l'Analyse à la Probabilité des Décisions Rendus à la Pluralité des Voix*. L'Imprimerie Royale, 1785.

M. Costantini, C. Groenland, and U. Endriss. Judgment aggregation under issue dependencies. In D. Schuurmans and M. P. Wellman, editors, *Proceedings of the 30th AAAI Conference on Artificial Intelligence*. Association for the Advancement of Artificial Intelligence, 2016.

F. De Cindio and S. Stortone. Experimenting LiquidFeedback for online deliberation in civic contexts. In M. A. Wimmer, E. Tambouris, and A. Macintosh, editors, *Proceedings of the 5th IFIP WG 8.5 International Conference on Electronic Participation (ePart)*, pages 147–158. Springer, 2013.

B. Debord. *Axiomatisation de Procédures d'Agrégation de Préférences*. PhD thesis, Université Scientifique et Médicale de Grenoble, 1987.

F. Dietrich and C. List. Arrow's theorem in judgment aggregation. *Social Choice and Welfare*, 29(1):19–33, 2007.

F. Dietrich and C. List. Judgment aggregation without full rationality. *Social Choice and Welfare*, 31:15–39, 2008.

C. L. Dodgson. *The Principles of Parliamentary Representation*. Harrison and Sons, 1884.

U. Endriss. Judgment aggregation. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 17. Cambridge University Press, 2016.

U. Endriss and R. De Haan. Complexity of the winner determination problem in judgment aggregation: Kemeny, Slater, Tideman, Young. In R. H. Bordini, E. Elkind, G. Weiss, and P. Yolum, editors, *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 117–126. International Foundation for Autonomous Agents and Multiagent Systems, 2015.

U. Endriss, R. de Haan, J. Lang, and M. Slavkovik. The complexity landscape of outcome determination in judgment aggregation. *Journal of Artificial Intelligence Research*, 69:687–731, 2020.

B. Escoffier, H. Gilbert, and A. Pass-Lanneau. The convergence of iterative delegations in liquid democracy in a social network. In D. Fotakis and E. Markakis, editors, *Proceedings of the 12th International Symposium on Algorithmic Game Theory (SAGT)*, pages 284–297. Springer, 2019.

P. Everaere, S. Konieczny, and P. Marquis. Counting votes for aggregating judgments. In A. Lomuscio, P. Scerri, A. Bazzan, and M. Huhns, editors, *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1177–1184. International Foundation for Autonomous Agents and Multiagent Systems, 2014.

B. Ford. Delegative democracy, 2002. URL `http://www.brynosaurus.com/deleg/deleg.pdf`. Unpublished manuscript [accessed: 2023-11-07].

U. Grandi and U. Endriss. Lifting integrity constraints in binary aggregation. *Artificial Intelligence*, 199:45–66, 2013.

J. Green-Armytage. Direct voting and proxy voting. *Constitutional Political Economy*, 26(2):190–220, 2015.

J. Harding. Proxy selection in transitive proxy voting. *Social Choice and Welfare*, 58(1):69–99, 2022.

A. Kahng, S. Mackenzie, and A. D. Procaccia. Liquid democracy: An algorithmic perspective. *Journal of Artificial Intelligence Research*, 70:1223–1252, 2021.

J. G. Kemeny. Mathematics without numbers. *Daedalus*, 88(4):577–591, 1959.

G. Köhler. *Choix Multicritère et Analyse Algébrique de Données Ordinales*. PhD thesis, Institut d'Informatique et de Mathématiques Appliquées de Grenoble, 1978.

L. A. Kornhauser and L. G. Sager. The one and the many: Adjudication in collegial courts. *California Law Review*, 81(1):1–59, 1993.

G. Kotsialou and L. Riley. Incentivising participation in liquid democracy with breadth-first delegation. In A. El Fallah Seghrouchni, G. Sukthankar, B. An, and N. Yorke-Smith, editors, *Proceedings of the 19th International Conference* *on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 638–644. International Foundation for Autonomous Agents and Multiagent Systems, 2020.

C. Lamboray. *Prudent Ranking Rules: Theoretical Contributions and Applications*. PhD thesis, Université Libre de Bruxelles and Université du Luxembourg, 2007.

C. Lamboray. A characterization of the prudent order preference function. *Mathematical Social Sciences*, 57(3):389–405, 2009a.

C. Lamboray. A prudent characterization of the Ranked Pairs Rule. *Social Choice and Welfare*, 32(1):129–155, 2009b.

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

C. List and P. Pettit. Aggregating sets of judgments: An impossibility result. *Economics & Philosophy*, 18(1):89–110, 2002.

A. Litvinenko. Social media and perspectives of liquid democracy: The example of political communication in the Pirate Party in Germany. In M. Gasco, editor, *Proceedings of the 12th European Conference on eGovernment (ECEG)*, pages 403–408. Academic Conferences, 2012.

P. Mancini. Why it is time to redesign our political system. *European View*, 14 (1):69–75, 2015.

J. C. Miller, III. A program for direct and proxy voting in the legislative process. *Public Choice*, 7(1):107–113, 1969.

P. Mongin. The doctrinal paradox, the discursive dilemma, and logical aggregation theory. *Theory and Decision*, 73(3):315–355, 2012.

K. Nehring and M. Pivato. Majority rule in the absence of a majority. *Journal of Economic Theory*, 183:213–257, 2019.

S. Neuland. Political awareness and symmetric vision evaluating the online government initiative Liquid Friesland from an openness perspective. *Maastricht Research Based Learning for Excellence*, 4:19–35, 2014.

D. Parkes and L. Xia. A complexity-of-strategic-behavior comparison between Schulze's rule and ranked pairs. In C. Sierra, editor, *Proceedings of the 26th AAAI Conference on Artificial Intelligence*, pages 1429–1435. Association for the Advancement of Artificial Intelligence, 2012.

K. Pearson. Note on regression and inheritance in the case of two parents. In *Proceedings of the Royal Society of London*, volume 58, pages 240–242. The Royal Society, 1895.

W. H. Riker. *Liberalism Against Populism: A Confrontation Between the Theory of Democracy and the Theory of Social Choice*. Waveland Press, 1988.

A. Saltelli. Making best use of model evaluations to compute sensitivity indices. *Computer Physics Communications*, 145(2):280–297, 2002.

M. Slavkovik. Not all judgment aggregation should be neutral. In A. Herzig and E. Lorini, editors, *Proceedings of the European Conference on Social Intelligence*, pages 198–211, 2014.

I. M. Sobol'. Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. *Mathematics and Computers in Simulation*, 55 (1–3):271–280, 2001.

G. Ten Broeke, G. Van Voorn, and A. Ligtenberg. Which sensitivity analysis method should I use for my agent-based model? *Journal of Artificial Societies and Social Simulation*, 19(1), 2016.

Z. Terzopoulou and U. Endriss. Neutrality and relative acceptability in judgment aggregation. *Social Choice and Welfare*, 55(1):25–49, 2020.

T. N. Tideman. Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4(3):185–206, 1987.

G. Tullock. *Toward a Mathematics of Politics*. University of Michigan Press, 1967.

M. Utke and U. Schmidt-Kraepelin. Anonymous and copy-robust delegations for liquid democracy. In A. Oh, T. Neumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Proceedings of the 36th Conference on Neural Information Processing Systems (NeurIPS)*, pages 69441–69463. Curran Associates, 2023.

T. M. Zavist and T. N. Tideman. Complete independence of clones in the ranked pairs rule. *Social Choice and Welfare*, 6(2):167–173, 1989.