Numerical Judgment Aggregation: Towards a General Framework

MSc Thesis *(Afstudeerscriptie)*

written by

Tisja Irene Smits
(born January 14th, 2000 in Hilversum, 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:
June 27, 2025

Members of the Thesis Committee:
Dr. Aybüke Özgün (chair)
Prof. Dr. Ulle Endriss (supervisor)
Prof. Dr. Davide Grossi
Dr. Ronald de Haan

INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

Abstract

This thesis introduces and develops the framework of Numerical Judgment Aggregation (NJA), extending classical judgment aggregation to settings where agents express numerical judgments over general variables. Unlike traditional models, NJA allows for more expressive inputs and outputs, accommodating real-world scenarios where uncertainty, gradation, or nuance play a central role.

Several classes of aggregation rules are defined and analyzed: calculation-based (e.g., mean), ordinal (e.g., median), and support-based (e.g., majority). The normative behavior of these rules is axiomatically evaluated by means of properties such as anonymity, unanimity, independence, autonomy, and monotonicity. A central result is the axiomatic characterization of support quota rules via the concept of winning coalitions, generalizing results from classical judgment aggregation.

To ensure well-structured outputs (e.g., as intervals or points), repair operations are introduced and studied in terms of their effects on axiomatic properties. The framework also accommodates the use of integrity constraints to formalize dependencies between issues, structural coherence, or domain-specific consistency requirements. The expressive power and flexibility of NJA are further illustrated through applications to preference aggregation, binary judgment aggregation, and societal trade-off modeling.

Overall, this thesis offers a unified and general approach to aggregating numerical judgments under structural constraints, providing both theoretical insight and practical relevance.

Acknowledgments

What a ride. This thesis, like most things in life, did not go as planned. In short, it began in enthusiasm, was derailed by disillusionment, hindered by perfectionism and guilt, revived by inspiration, derailed again by existential dread, and then—in cautious optimism—found its way to completion, one day at a time. So, before I thank those who helped me through this turbulent phase of my life, let me first thank the thesis itself—for waiting until I was ready. Thanks, I guess.

First and foremost, to my supervisor, Ulle, thank you for your trust, patience, and support throughout this project, especially regarding my many a fruitless efforts and stubbornness to find some big, groundbreaking result. Thank you for being willing to supervise a project that took much longer than either of us expected, and most importantly, thank you for letting me find my way back to this project on my own terms.

To Computational Social Choice and Game Theory, to Ulle for teaching these courses, and to my fellow students with whom I did the group project that turned into this thesis—thank you for sparking (and re-sparking) my inspiration for this project.

To Floris, my love, my co-thesis-survivor, thank you for bringing me lots and lots of cookies and tea, for giving me daily structure, not letting me work too long or too late and making me break for lunch, for making space for me—literally and emotionally—and for making life a little softer in the process. We joked about it being a race, but now we have finally reached the finish line. (Still, for the record, I won.)

To Nicole, my geographically challenged bestie, thank you for our effortless friendship. Thank you for our snacking and surfing adventures, our travels through Mexico, my stay in Switzerland, teaching me to snowboard, saving my life in Ecuador, making me laugh until my stomach hurts, and for making the world feel smaller and lighter. I hope to see you soon, wherever that may be.

To my other amazing friends. To Tim, for all the maymays and Overcooked cucumber shenanigans. To Jaime, for your silent friendship. The past years have been hard, but even if we were not in sync, I knew you were there. To the infamous Thuisfront, for the well-intentioned tough love from time to time, and helping me stay grounded. A special thanks to the 'lustrumcie' for the very poorly timed yet 'gezellige' trip to Bologna. To Norah, for Italian carbs and curly girl frustrations. To Liotte, for thesis misery shared. To Manon, Nadèche, and Madelon, my former housemates, for the place I called home, for letting me be, but also for checking up on me every once in a while.

To Touchée and Ferocious, for freeing my body when my mind felt like a cage, for the music, the movement, the confidence, the support and the fierce energy.

To Kayo, for teaching me when to step on the gas and when to break—literally.

To my parents, for your continuous support and sincere curiosity. Thanks, Mom, for regularly checking up on me and always knowing when something is wrong, and for bringing me churros in 'de laatste loodjes.'

A heartfelt thanks to everyone—friends, family, colleagues, and innocent bystanders—who pretended to understand and patiently listened to me trying to explain the subject of my thesis in 'jip-en-janneketaal'. Thank you for teaching me a valuable lesson; explaining your thesis to someone is often harder than proving the theorem itself.

There is no neat moral to this story—only that perseverance looks different for everyone. For me, it looked like striving, dancing, crying, giving up, starting again, more dancing, and eventually learning to be okay with good enough. I am not there yet. But I am not perfect. And neither is this thesis. And that is okay.

So finally, thanks to me.

Contents

Chapter 1

Introduction

Social choice theory is a foundational field that branches out into economics, political science, philosophy, logic, and theoretical computer science. It studies how individual preferences, judgments, or welfare can be aggregated to form collective decisions. Originating from the work of pioneers such as Nicolas de Condorcet (1785), Kenneth Arrow (1951), and Amartya Sen (1970), social choice theory seeks to provide formal models for analyzing collective decision-making in diverse contexts, including electoral systems, public policymaking, and multi-agent systems. At its core, social choice theory addresses fundamental questions such as: How can individual inputs be combined to form a group decision? What are the properties of different voting systems? What properties should an aggregation procedure satisfy to ensure rationality, consistency, and fairness? What limitations, impossibilities, or paradoxes arise when aggregating individual inputs?

Early work in social choice focuses on preference aggregation, where individuals' preference orderings of alternatives (i.e., rankings) are combined into a single, collective ordering or choice over these alternatives. One of the most influential results in social choice theory is Arrow's Impossibility Theorem (1950), which states that no preference aggregation system satisfies five plausible fairness axioms. Almost two centuries earlier, Condorcet had already discovered a special case of what would later be Arrow's Theorem. The Condorcet paradox (1785), as it is called, shows that collective preferences may become cyclic, even if individual preferences are transitive (see Figure 1.1 below). The result implies that it is logically impossible for any preference aggregation system to guarantee that a winner will have support from a majority of agents.

agent 1	$a \succ b \succ c$
agent 2	$b \succ c \succ a$
agent 3	$c \succ a \succ b$
majority	$a \succ b \succ c \succ a \succ \dots$

Figure 1.1: Condorcet paradox

A much younger discipline within social choice is judgment aggregation. Whereas much of preference aggregation focuses on ordinal judgments over a set of alternatives, standard judgment aggregation (JA) deals with binary judgments (i.e., yes/no decisions) on logically interconnected propositions. In addition to fairness, JA is mainly concerned with logical consistency. The formal work on judgment aggregation stems from the "doctrinal paradox" in the jurisprudence literature of the late twentieth century (e.g., Kornhauser & Sager, 1986). The paradox illustrates how aggregating individually consistent judgments with the majority rule can result in self-contradictory collective judgments (see Figure 1.2 below).

	p	q	$p \wedge q$
agent 1	yes	no	no
agent 2	no	yes	no
agent 3	yes	yes	yes
majority	yes	yes	no

Figure 1.2: Doctrinal paradox

Although much of social choice theory focuses on qualitative judgments (e.g., voting for one candidate over another or accepting/rejecting a proposition), many real-world decision-making problems involve quantitative judgments. For example, in forecasting, policymaking, and risk assessment, individuals express their judgments in numerical form, such as probabilities, percentages, prices, or ratings. These numerical judgments can take different forms (e.g., points, intervals, or arbitrary sets), representing uncertainty, suspension of judgment, or underdetermination.

Aggregating numerical judgments is not completely unknown, even within JA. For example, Pauly and van Hees (2006) present a framework that generalizes binary judgment aggregation by incorporating many-valued logic to express degrees of acceptance and rejection, as opposed to a binary notion when using classical logic. The theory of probabilistic opinion pooling (see e.g., Elkin & Pettigrew, 2025) goes even further by aggregating judgments in a probabilistic setting as opposed to a logical one. In opinion pooling, judgments are interpreted as subjective probabilities, called credences, instead of degrees of truth or acceptance. Whereas JA is based on logic and formal aggregation procedures and focuses on ensuring logical consistency, opinion pooling is based on probability theory and uses statistical measures such as weighted averaging methods, and it focuses on ensuring probabilistic coherence.

Although JA and probabilistic opinion pooling are concerned with different sets of issues, that is, different agendas, both frameworks are contingent on the algebraic structure of the agenda, whether logically or probabilistically. Dietrich and List (2017) explore the idea of opinion pooling for general agendas. However, they lose the semantics of the agenda in the process, and hence expressivity. To preserve the semantics, we can incorporate so-called integrity constraints that ensure agent rationality and outcome feasibility without imposing an internal structure on the agenda. Most of the existing work (e.g., Endriss, 2018; Grandi & Endriss, 2010) encodes these constraints as propositional formulas. Here, we extend this to mathematical formulas of equality and inequality.

This thesis develops and explores a formal framework for numerical judgment aggregation (NJA) that incorporates general agendas in combination with integrity constraints, making it both expressive and versatile. We aim to formalize the concepts of numerical judgments and define their structural properties, define and evaluate (classes of) judgment aggregation rules, explore their axiomatic properties, introduce pragmatic methods for repairing undesired aggregation outputs, and demonstrate several applications of NJA, such as the simulation of binary JA.

The structure of the thesis is as follows. Chapter 2 introduces the framework of NJA. The general model is presented in Section 2.1, formalizing the concept of numerical judgments and their properties. Section 2.2 introduces different (classes of) judgment aggregation rules. The chapter concludes with Section 2.3, which discusses pragmatic operations to repair the format of unwanted output. Chapter 3 discusses the properties of aggregation rules. These properties are formalized into axioms in Section 3.1, after which Section 3.2 analyzes which rules satisfy which properties, characterizes the support-quota rules, and discusses how the use of repair operations affects rule properties. Chapter 4 introduces integrity constraints for NJA and discusses their pragmatic importance in Section 4.1.

Section 4.2 defines the formal notions, and Sections 4.3, 4.4, and 4.5 provide some results on the consistency of judgments with respect to certain constraints. Chapter 5 covers several applications of the framework, such as the embedding of preference aggregation in Section 5.1, the simulation of binary JA in Section 5.2, and a point-valued model for the analysis of societal trade-offs in Section 5.3. Chapter 6 concludes the thesis with a summary, discussion, and outline of potential further research. Lastly, Appendix A provides some notation and definitions of mathematical concepts that are used throughout this thesis.

Chapter 2

Numerical Judgment Aggregation

The first step in developing a general judgment aggregation framework is to change the content of the agenda. Instead of propositions (as for JA) or probabilities (as for probabilistic opinion pooling), agenda-items have no presupposed meaning. Instead, the context imposes their meaning. Agenda-items might still constitute logical formulas that require a degree of truth, but they could also represent other objects, such as the number of tiles flipped in a municipality during the 'NK Tegelwippen',[1] the probability that it will rain tomorrow or the amount of funding a project should receive. Agenda-items can represent any real-world or mathematical object that can be ascribed a number.

This new type of agenda-items requires a different format for judgments. It is no longer a question of yes or no, so we expand the response options to the set of real numbers. This allows for nuance in agents' judgments. Furthermore, agents are allowed to report a set of numbers instead of a single one. In this way, agents can express indifference or doubt by reporting multiple numbers, or they can withhold judgment on certain issues by reporting the empty set. In the current version of the framework, we only consider closed intervals to be proper input, but this can be extended to arbitrary sets of real numbers with some modifications to certain rules or axioms.

This chapter defines the basic framework and introduces different classes of aggregation rules, including some based on common procedures such as weighted aggregation and majority voting. In the last section, we discuss the use of pragmatic operations that repair possibly undesired output.

2.1 The Framework

Let $N = \{1, \ldots, n\}$ with $n \geq 2$ be a finite set of agents concerned with reaching a collective judgment on certain issues. These issues are represented by a finite (nonempty) set of variables, $X = \{x, y, \ldots\}$, referred to as the *agenda*. Each agenda-item $x \in X$ requires a so-called *judgment value* in the form of a (possibly empty) set of real numbers; $J(x) \subseteq \mathbb{R}$. This value represents the quantitative worth of the variable. A *judgment function* J, or *judgment* for short, maps every agenda-item $x \in X$ to such a judgment value; $J : X \to 2^{\mathbb{R}}$. In this way, the term 'judgment value' captures both the mathematical sense of the word (i.e., value of the function) and the economic sense of the word (i.e., value to an agent).

Definition 2.1 (Judgment function properties)**.** A judgment function J on an agenda X is called

[1] National championship in tile flipping. The goal of the competition is to replace as many tiles as possible with green. See https://www.nk-tegelwippen.nl/

- *congruous*[2] if for no $x \in X$, $J(x) = \varnothing$. That is, reporting no possible values is inappropriate.
- *trivial* if for some $x \in X$, $J(x) = \mathbb{R}$, in the sense that it is trivial to report that a variable can take any possible value. It is *nontrivial* if there is no such $x \in X$.
- *ranged* if for all $x \in X$, $J(x) = [a, b]$ for some $a, b \in \mathbb{R}$. That is, the judgment consists of closed intervals. The intended reading of $J_i(x) = [a, b]$ is that agent i judges the value of x to be somewhere between a and b.
- *point-valued* if for all $x \in X$, $J(x) = \{k\}$ for some $k \in \mathbb{R}$. That is, the judgment assigns singletons to variables, representing a precise numerical opinion.
- *bivalent* if for all $x \in X$, $J(x) = \{k\}$ for some $k \in \{0, 1\}$.

Note that bivalent judgments are, by definition, point-valued; point-valued judgments are ranged because a singleton $\{k\}$ can be denoted as the interval $[k, k]$; and ranged judgments are both congruous and nontrivial. For better readability, we will often use the notation $\boldsymbol{k}$ to denote a singleton value $\{k\}$; and when J is clear from context, we use the shorthands a_i^x and b_i^x, for every $x \in X$ and every $i \in N$, to refer to a_i^x and b_i^x such that $J_i(x) = [a_i^x, b_i^x]$.[3]

Let $\mathcal{J}(X)$ be the set of all judgment functions on the agenda X. Agents 1 through n are each required to report a ranged judgment function. These are collected into a so-called *profile* $\boldsymbol{J} = (J_1, \dots, J_n)$. A *judgment aggregation rule* $F : \mathcal{J}(X)^n \to \mathcal{J}(X)$ is a function that maps a profile of individual judgment functions to a single collective judgment function $F(\boldsymbol{J})$.

In general, we are interested in finding aggregation rules that preserve the notions of congruity and nontriviality. Whether the output should also be ranged or point-valued, for example, depends on the application. However, in the standard model, we always assume that the input is ranged, unless explicitly stated otherwise.

2.2 Aggregation Rules

This section is dedicated to the different (classes of) aggregation rules: calculation-based rules, ordinal rules, and support-based rules. We distinguish these classes based on how the input profile is processed. If the collective judgment is determined directly by some calculation on the judgment values in the profile, we speak of a calculation-based rule. Ordinal rules determine the collective judgment based on the order of the values from the individual judgments. And if a rule determines which values make it into the collective judgment based on how much support they have from the individual agents, it is called support-based. Note that these classes are not defined, but rather a natural way to classify the rules presented here.

2.2.1 Calculation-Based Rules

As mentioned, calculation-based rules determine the collective judgment directly through some (equation-based) calculation on the judgment values in the profile. In the case of a point-valued profile, a simple example of such a rule is one that returns the sum of all points. But the calculation could also be more complicated, involving, for example, logarithmic or exponential functions, or weights assigned to agents' judgments. More

[2]Meaning 'conforming to the circumstances or requirements of a situation; appropriate' (Merriam-Webster, n.d.).

[3]For (half)-open intervals, or even arbitrary sets, we would assign $a_i^x = \inf J_i(x)$ and $b_i^x = \sup J_i(x)$.

practical rules are those that use some method to average the individual judgments. There are many different averaging techniques. The most common is the arithmetic (or simple) mean, which is used in the mean rule below.

The Mean Rule

For a variable $x \in X$ and a profile $\boldsymbol{J} \in \mathcal{J}(X)^n$, the collective judgment value of x is an interval from the mean infimum until the supremum:

$$\begin{aligned} F_{\text{mean}}(\boldsymbol{J})(x) &:= \frac{1}{n}\sum_{i=1}^{n} J_i(x) \\ &= \left[\frac{1}{n}\sum_{i=1}^{n} a_i^x, \frac{1}{n}\sum_{i=1}^{n} b_i^x\right] \end{aligned} \tag{2.1}$$

Weighted Mean Rules

In the (simple) mean rule above, each agent's judgment carries equal weight. A *weighted mean rule* $F_{\text{mean},w}$, on the other hand, is defined by a *weight function* $w : N \to [0,1]$ with $\sum_{i=1}^{n} w_i = 1$, assigning each agent a (normalized) weight. These weights represent the voting power of the agents in the model. They can be used to represent expertise, trustworthiness, participation frequency, etc. The simple mean rule is a special case where all agents have equal voting power, i.e., $w_i = \frac{1}{n}$ for all $i \in N$. Another special case is when there is a single agent $i \in N$ for which $w_i = 1$. We call this a *dictatorship* because there is an agent with the power to single-handedly determine the outcome.

For a variable $x \in X$ and a profile $\boldsymbol{J} \in \mathcal{J}(X)^n$, the collective judgment value of x is an interval from the weighted mean infimum until the weighted mean supremum:

$$\begin{aligned} F_{\text{mean},w}(\boldsymbol{J})(x) &:= \sum_{i=1}^{n} w_i \cdot J_i(x) \\ &= \left[\sum_{i=1}^{n} w_i \cdot a_i^x, \sum_{i=1}^{n} w_i \cdot b_i^x\right] \end{aligned} \tag{2.2}$$

We can extend this approach using weights even further by changing w into a function on the set of variables, $w : X \to [0,1]^N$. It could be used to specify the agents' voting power on certain topics, e.g., their area of expertise.

2.2.2 Ordinal Rules

An ordinal is a number that represents a position or rank in a sequential order. Ordinal rules determine the collective judgment based on the order of the values from the individual judgments. An example is a rule that takes the individual judgment value with the smallest supremum.

The Median Rule

Definition 2.2 (Median function)**.** The (upper) median is a function that takes a set S of $|S| = m$ numbers, orders it from smallest to greatest, and returns the middle value:

$$\text{med}(S) := S_{\lceil \frac{m+1}{2} \rceil}$$

For a variable $x \in X$ and a profile $\boldsymbol{J} \in \mathcal{J}(X)^n$, the collective judgment value of x is an interval from the median infimum until the median supremum:

$$F_{\text{med}}(\boldsymbol{J})(x) := \left[\text{med}\left\{a_1^x, \ldots, a_n^x\right\}, \text{med}\left\{b_1^x, \ldots, b_n^x\right\}\right] \tag{2.3}$$

The Min and Max Rules

For a variable $x \in X$ and a profile $\boldsymbol{J} \in \mathcal{J}(X)^n$, the collective judgment value of x is an interval from the minimum or maximum infimum until the minimum or maximum supremum, respectively:

$$F_{\min}(\boldsymbol{J})(x) := \left[\min_{i\in N} a_i^x, \min_{i\in N} b_i^x\right] \tag{2.4}$$

$$F_{\max}(\boldsymbol{J})(x) := \left[\max_{i\in N} a_i^x, \max_{i\in N} b_i^x\right] \tag{2.5}$$

Bound-Quota Rules

A bound-quota rule F_{bq} is defined by a function $q : X \to \{0, 1, \dots, n+1\}$. F_{bq} is called *uniform* if q maps all variables to the same number λ.

For a variable $x \in X$ and a profile $\boldsymbol{J} \in \mathcal{J}(X)^n$, let

$$a_q := \min\{a : |\{i \in N : a_i^x \leq a\}| \geq q(x)\}$$
$$b_q := \min\{b : |\{i \in N : b_i^x \leq b\}| \geq q(x)\}$$

The collective judgment value of x is the interval from the $q(x)$-th infimum until the $q(x)$-th supremum:

$$F_{\text{bq}}(\boldsymbol{J})(x) := [a_q, b_q] \tag{2.6}$$

The (strict) majority bound rule is defined by $\lambda = \lceil \frac{n+1}{2} \rceil$. It returns the interval with the minimum infimum and supremum for which a majority report that the bounds should be at least as small. Intuitively, a bound-quota rule seems to order the bounds from smallest to greatest and then returns the $q(x)$-th bound. The upper median is the $\lceil \frac{n+1}{2} \rceil$-th value. So, the majority bound rule is, in fact, equivalent to the median rule (2.3) with $\lambda = \lceil \frac{n+1}{2} \rceil$. Similarly, the min and max rules, 2.4 and 2.5, can be classified as bound-quota rules, with $\lambda = 1$ and $\lambda = n$, respectively.

Other Versions of the Bound-Quota Rules

The original definition approaches the $q(x)$-th infimum from below, i.e., as if the infimums were ordered from smallest to greatest:

$$a_q := \min\{a : |\{i \in N : a_i^x \leq a\}| \geq q(x)\} \tag{option A}$$

However, we could also approach the $q(x)$-th infimum from above:

$$a_q := \max\{a : |\{i \in N : a_i^x \geq a\}| \geq q(x)\} \tag{option B}$$

Notice how 'min' and '$\leq$' in option A have respectively changed to 'max' and '$\geq$' in option B.

In the current definition, we opt for option A for both the infimum and the supremum, to show the most intuitive connection to the median ($\lambda = \lceil \frac{n+1}{2} \rceil$), min ($\lambda = 1$) and max ($\lambda = n$) rules. If we use option B for both the infimum and the supremum, we get different results because we approach the bounds from the other direction. The rule with $\lambda = 1$ is then equivalent to the max rule instead, and the rule with $\lambda = n$ to the min rule, so precisely the other way around. Furthermore, the rule with $\lambda = \lceil \frac{n+1}{2} \rceil$ would be the lower median rule instead of the (upper) median rule.

If we do not use the same option for the infimum and supremum, we get even more deviant results. For example, if we use option A for the infimum and option B for the supremum, we get the interval enclosure of the nomination rule (2.7) with $\lambda = 1$ and the interval enclosure of the unanimity rule (2.8) with $\lambda = n$. These rules belong to the class of support-based rules, which will be discussed in the next subsection, but their interval enclosures are technically ordinal rules with this alternate definition. It is left to the reader to check what happens when we use option B for the infimum and option A for the supremum.

2.2.3 Support-Based Rules

Support-based rules determine which values make it into the collective judgment based on how much support they have from the agents. We formally define the notion of support as follows:

Definition 2.3 (Support)**.** For a profile $\boldsymbol{J}$, a variable $x \in X$, and a number $k \in \mathbb{R}$, the support $N^{\boldsymbol{J}}_{x,k} \subseteq N$ is a set that consists of agents that consider k a possible value for x:

$$N^{\boldsymbol{J}}_{x,k} := \{i \in N : k \in J_i(x)\}$$

The *amount* of support $\#N^{\boldsymbol{J}}_{x,k}$ is then given by $|\{i \in N : k \in J_i(x)\}|$.

Support-based rules generally do not guarantee ranged output. Depending on the specific input profile, the collective might be a union of intervals rather than a single closed interval.

The Nomination Rule

For a variable $x \in X$ and a profile $\boldsymbol{J} \in \mathcal{J}(X)^n$, the collective judgment value of x is the collection of all reported values:

$$F_{\cup}(\boldsymbol{J})(x) := \bigcup_{i \in N} J_i(x) \tag{2.7}$$

If the reported values overlap in such a way that their union is again a closed interval, we have $F_{\cup}(\boldsymbol{J})(x) = [\min_{i \in N} a_i^x, \max_{i \in N} b_i^x]$.

The Unanimity Rule

The unanimity rule returns only the values that everyone considers possible. For a variable $x \in X$ and a profile $\boldsymbol{J} \in \mathcal{J}(X)^n$, the collective judgment value of x is the intersection of all reported values:

$$F_{\cap}(\boldsymbol{J})(x) := \bigcap_{i \in N} J_i(x) \tag{2.8}$$

Given an input profile of ranged judgment functions, this rule always returns a (possibly empty) interval, namely the interval from the maximum infimum until the minimum supremum. So, for a ranged profile, we have:

$$F_{\cap}(\boldsymbol{J})(x) = \begin{cases} \left[\max\limits_{i \in N} a_i^x, \min\limits_{i \in N} b_i^x\right] & \text{if } \max_{i \in N} a_i^x \leq \min_{i \in N} b_i^x \\ \varnothing & \text{otherwise} \end{cases}$$

The Mode Rule

For a variable $x \in X$ and a profile $\boldsymbol{J} \in \mathcal{J}(X)^n$, the collective judgment value of x contains the values k that receive the greatest amount of support:

$$F_{\text{mode}}(\boldsymbol{J})(x) := \arg\max_{k \in \mathbb{R}} \#N^{\boldsymbol{J}}_{x,k} \tag{2.9}$$

The mode rule does not guarantee an interval outcome. For example, consider $n = 3$ with $J_1(x) = [0, 2]$, $J_2(x) = [3, 5]$, and $J_3(x) = [1, 6]$. We get $F_{\text{mode}}(\boldsymbol{J})(x) = [1, 2] \cup [3, 5]$.

Support-Quota Rules

Similarly to a bound-quota rule (2.6), a *support-quota rule* F_{sq} is defined by a function $q : X \to \{0, 1, \dots, n+1\}$. F_{sq} is called *uniform* if q maps all variables to the same number λ.

For a variable $x \in X$ and a profile $\boldsymbol{J} \in \mathcal{J}(X)^n$, the collective judgment value of x is the set of numbers that have the required amount of support, i.e., that are included in the reported intervals of at least $q(x)$-many agents:

$$F_{\text{sq}}(\boldsymbol{J})(x) := \left\{k : \#N^{\boldsymbol{J}}_{x,k} \geq q(x)\right\} \tag{2.10}$$

The (strict) majority support rule is defined by $\lambda = \lceil \frac{n+1}{2} \rceil$. It returns the set of numbers that are supported by a strict majority of agents. In binary JA, the majority support rule and the median rule (2.3) are equivalent. Here, however, they are not. Consider, for example, a point-valued profile, where all agents report different values. The majority support rule will then return the empty set because no number has the required amount of support. The median rule, however, does not depend on the amount of support for values, so it will just return the interval with the median bounds. For bivalent profiles, we do have an equivalence between the majority support rule and the median rule (and hence the majority bound rule (2.6)).

In a way, all support-based rules defined here can be classified as support-quota rules. The nomination rule (2.7) can be defined with $\lambda = 1$, and the unanimity rule (2.8) can be defined with $\lambda = n$. Strictly speaking, the mode rule (2.9) is not a support-quota rule, because the function $q(x)$ is defined independently of the input profile. However, intuitively, the mode rule can be described by the greatest $q(x)$ for every $x \in X$ such that $F_{\text{sq}}(x)$ is nonempty.

2.3 Repair Operations

Many support-based rules do not necessarily return collective judgment values in the form of intervals, but arbitrary sets of points instead.[4] For many applications of judgment aggregation or collective decision-making in general, it is reasonable to believe that the outcome is a means to an end; for example, for policymaking. Some applications might not be suitable for dealing with such arbitrary sets. They require more manageable output. If a desired rule does not guarantee such output, we can apply a so-called *repair operation*. Repair operations are functions that force the outcome into a specific format.

Point Repair

If the context requires a point-valued outcome, arbitrary sets and intervals are not precise enough. To solve this, we can apply an operation that turns a set into a single point. A point repair is basically a tie-breaking rule applied to the outcome.

A straightforward operation is to take the mean. However, depending on the aggregation rule, agents might be able to manipulate the outcome by reporting a different judgment than their true one. Furthermore, the mean is only defined for finite sets. An alternative is to take a random point. Although randomness is not manipulable, a random point might not sufficiently represent the collective. Another option is to take the median of a (finite) set, or the midpoint in case of an interval:

Definition 2.4 (Midpoint function)**.** The midpoint of an interval with infimum a and supremum b is $(a + b)/2$.

[4] That is, arbitrary in their structure, not in the way the points are obtained.

Interval Repair

For turning arbitrary sets into intervals, we take inspiration from a concept in geometry called the ‘convex hull’ (Fan, 1959). The convex hull operator is a closure operator that enwraps a set of points as if you would put a rubber band around it. Formally, a set of points is *convex* if it contains every line segment between two points in the set. The *convex hull* of a set is then defined as the smallest convex set that contains it. Within our one-dimensional domain of real numbers, a convex set is simply an interval. Consequently, the convex hull of a set of numbers is its interval enclosure. The convex hull of a compact set of numbers is a closed interval, namely the interval from the minimum to the maximum. If the original set is open, then the convex hull is also open. For example, the convex hull of all positive even numbers is the half-open interval $[2, \infty)$.

Definition 2.5 (Convex hull operation)**.** Given a judgment value $J(x)$ for some variable $x \in X$, the convex hull operation returns the smallest interval I that contains all the numbers in $J(x)$:

$$\text{hull}\,(J(x)) = \underset{I \supseteq J(x)}{\arg\min}\,|I|$$

Another way to turn an arbitrary set into an interval is to find a convex set inside it rather than around it. We call the largest such convex set the *convex core*. Think of this as the set being sieved with an increasingly finer sieve. The first residue we find is the convex core. Note that there might be several convex cores. We could apply a context-dependent tie-breaking rule to determine a unique convex core.

Definition 2.6 (Convex core operation)**.** Given a judgment value $J(x)$ for some variable $x \in X$, the convex core operation returns the largest interval I that is contained in $J(x)$:

$$\text{core}\,(J(x)) \in \underset{I \subseteq J(x)}{\arg\max}\,|I|$$

The convex hull operation and the convex core operation are both particularly suitable for support-based rules, such as the mode rule (2.9) and the nomination rule (2.7), because these do not always return intervals.

Chapter 3

Properties of Aggregation Rules

The choice of aggregation rule depends on the purpose and context. However, to know whether a rule is appropriate, we have to know how it behaves, i.e., what properties it does or does not have. Fairness, accuracy, and consistency are considered the most important properties to have, but their interpretations vary widely. Fairness ensures that all agents are treated equitably. For some applications, that means that each agent's judgment is treated equally, whereas for others, it means that an agent's say in the matter is determined by their authority or expertise. Accuracy describes how well the aggregated judgment reflects the true collective judgment. And consistency requires that the result remain coherent between variables or profiles, or under certain conditions, or that judgments adhere to the rules of, for example, logic or probability.

In this chapter, we use the axiomatic method to study the properties of judgment aggregation rules. The first section introduces some normatively appealing properties related to fairness, accuracy, and consistency, and provides formal definitions, called axioms. The second section analyzes the aggregation rules and repair operations presented in the previous chapter by means of these axioms and gives some characterizations of the support-quota rules. In the last section, we discuss what integrity constraints are and how to implement them in NJA.

3.1 Axioms

This section presents some properties of aggregation rules, formally defined as axioms. Some originate directly from standard judgment aggregation, some are adapted for numerical judgments, and some are completely new.

Definition 3.1 (Anonymity)**.** A rule F is *anonymous* if all agents are treated equally. For any profile $\boldsymbol{J} \in \mathcal{J}(X)^n$ and permutation $\pi : N \to N$:

$$F(J_1, \dots, J_n) = F(J_{\pi(1)}, \dots, J_{\pi(n)})$$

The axiom states that the collective judgment should be the same, no matter how the agents might be interchanged. For example, weighted mean rules (2.2) are not anonymous, unless the weights are uniformly distributed.

Anonymity is an axiom that ensures a certain degree of fairness. However, this might be at the expense of accuracy. For example, if some agents have more expertise in a certain subject than others, we would acknowledge their expertise by relaxing the axiom.

Definition 3.2 (Dictatoriality)**.** F is *dictatorial* if there exists an agent $i \in N$ such that for any profile $\boldsymbol{J} \in \mathcal{J}(X)^n$:

$$F(\boldsymbol{J}) = J_i$$

Otherwise, F is *nondictatorial*. If F is dictatorial, it is called a *dictatorship*.

A dictatorship is a rule where an agent has the power to single-handedly determine the outcome. The collective judgment is a direct copy of their individual judgment. As previously mentioned, a weighted mean rule (2.2) is a dictatorship for the special case where $w_i = 1$ for some agent $i \in N$.

It is controversial to say that a dictatorship is fair in terms of representation and democratic voting rights. However, by taking away agents' voting power, it is also not possible for them to manipulate the outcome in their favor. Moreover, when assuming that all individual agents are rational and consistent, a dictatorship guarantees that the collective outcome is rational and consistent as well (Dietrich & List, 2007b).

Definition 3.3 (Unanimity)**.** A rule is *(globally) unanimous* if, whenever all agents report the same judgment, it is also the collective judgment. For any profile $\boldsymbol{J} \in \mathcal{J}(X)^n$:

$$\text{if } J_1 = \cdots = J_n, \text{ then } F(\boldsymbol{J}) = J_1 = \cdots = J_n$$

A rule is *locally unanimous* if, whenever all agents report the same judgment value, it is also the collective judgment value. For any variable $x \in X$ and profile $\boldsymbol{J} \in \mathcal{J}(X)^n$:

$$\text{if } J_1(x) = \cdots = J_n(x), \text{ then } F(\boldsymbol{J})(x) = J_1(x) = \cdots = J_n(x)$$

Local unanimity is strictly stronger than global unanimity. It is obvious that the former implies the latter. If agents agree with each other, the unanimity axioms require that such agreement be represented in the collective. Agreement is often hard to come by; these axioms state that the rule should take advantage of it when possible. In some situations, it may instead make sense to relax the axioms. For example, individuals have different private evidence to justify their judgment, but when combined, the full evidence is inconsistent.

Definition 3.4 (Unanimous Support and Refutation)**.** A rule F is *unanimously supporting* if, whenever all agents include the same number in their judgment value, this number will also be in the collective judgment value. For any variable $x \in X$ and profile $\boldsymbol{J} \in \mathcal{J}(X)^n$:

$$J_1(x) \cap \cdots \cap J_n(x) \subseteq F(\boldsymbol{J})(x)$$

A rule F is *unanimously refuting* if, whenever all agents exclude the same number in their judgment value, this number will also not be in the collective judgment value. For any variable $x \in X$ and profile $\boldsymbol{J} \in \mathcal{J}(X)^n$:

$$F(\boldsymbol{J})(x) \subseteq J_1(x) \cup \cdots \cup J_n(x)$$

Unanimous support and unanimous refutation together imply local unanimity and thus global unanimity. However, separately, there are no direct implications between these properties. A rule that is locally unanimous, but not unanimously refuting, is the mean rule (2.1). A rule that is locally unanimous, but not unanimously supporting, is the rule that returns the judgment value for a variable on which all agents agree and otherwise returns the empty set.

Definition 3.5 (Neutrality)**.** A rule F is *(globally) neutral* if all issues are treated equally at the level of judgment values. For any variables $x, y \in X$ and profile $\boldsymbol{J} \in \mathcal{J}(X)^n$:

$$\text{if } J_i(x) = J_i(y) \text{ for all } i \in N, \text{ then } F(\boldsymbol{J})(x) = F(\boldsymbol{J})(y)$$

A rule F is *locally neutral* if all issues are treated equally at the level of individual numbers. For any variables $x, y \in X$, profile $\boldsymbol{J} \in \mathcal{J}(X)^n$ and $k \in \mathbb{R}$:

$$\text{if } k \in J_i(x) \Leftrightarrow k \in J_i(y) \text{ for all } i \in N, \text{ then } k \in F(\boldsymbol{J})(x) \Leftrightarrow k \in F(\boldsymbol{J})(y)$$

Global neutrality states that if all judgments are the same for two issues individually, then they should also be so collectively. Local neutrality is a stronger axiom in the sense that neutrality should also hold for individual numbers. That is, if agents agree that a certain number should be in or out for both variables, then it should be in both or neither of the variables' collective judgment values as well. It is left to the reader to check that local neutrality indeed implies global neutrality. Sometimes, issues might not be equally important or urgent. In that case, we would relax neutrality.

Definition 3.6 (Independence)**.** A rule F is *(globally) independent* if it decides on one issue at a time at the level of judgment values. For any variable $x \in X$ and profiles $\boldsymbol{J}, \boldsymbol{J'} \in \mathcal{J}(X)^n$:

$$\text{if } J_i(x) = J'_i(x) \text{ for all } i \in N, \text{ then } F(\boldsymbol{J})(x) = F(\boldsymbol{J'})(x)$$

A rule F is *locally independent* if it decides on one issue at a time at the level of individual numbers. For any variable $x \in X$, profiles $\boldsymbol{J}, \boldsymbol{J'} \in \mathcal{J}(X)^n$ and $k \in \mathbb{R}$:

$$\text{if } k \in J_i(x) \Leftrightarrow k \in J'_i(x) \text{ for all } i \in N, \text{ then } k \in F(\boldsymbol{J})(x) \Leftrightarrow k \in F(\boldsymbol{J'})(x)$$

The independence axioms look similar to the neutrality axioms. The difference is that neutrality states something about different variables within the same profile, and independence states it about the same agenda-item for different profiles. Global independence states that if each agent reports the same judgment for an agenda-item in one profile as in another, then their collective judgments should be the same for both profiles. Local independence is a stronger axiom in the sense that this independence should also hold for individual numbers. It is left to the reader to check that local independence indeed implies global independence.

The idea behind the independence axiom is that the collective judgment of an issue should only rely on its individual judgments, not on any judgments on other issues. For example, how many die rolls it takes to roll a six and the average global surface temperature do not depend on each other, so their collective judgments should not either. However, sometimes it might not be so obvious or even rational to consider all issues independently. The global surface temperature might not influence the number of die rolls, but it does influence the total glacier mass and the average sea level.

Definition 3.7 (Autonomy)**.** A rule F is *autonomous* if all numbers are treated equally, regardless of their numerical nature. For any variable $x \in X$, profile $\boldsymbol{J} \in \mathcal{J}(X)^n$ and $k, k' \in \mathbb{R}$:

$$\text{if } k \in J_i(x) \Leftrightarrow k' \in J_i(x) \text{ for all } i \in N, \text{ then } k \in F(\boldsymbol{J})(x) \Leftrightarrow k' \in F(\boldsymbol{J})(x)$$

The autonomy axiom and the local neutrality and independence axioms look similar, but autonomy is a form of neutrality specifically for numbers, not agenda-items. That is, autonomous rules have a certain indifference toward the numerical nature of numbers. They do not see numbers as positive or negative, odd or even, or larger or smaller than other numbers; they are just objects. That is why calculation-based and ordinal rules are not autonomous.

An example of an autonomous rule is a support-based rule like the unanimity rule (2.8). According to that rule, if all agents agree on whether two numbers should be both in or out, then either they are both reported by every agent, and thus in the collective judgment, or both are not reported by any agent, and thus not in the collective judgment. Note that the axiom itself does not require that both numbers be in the collective judgment if they are unanimously reported; it only requires that the rule treat them symmetrically.

Definition 3.8 (Monotonic Support)**.** A rule F has *monotonic support* if any collectively accepted value is still accepted if it receives additional support. For any agent $i \in N$, variable $x \in X$, profile $\boldsymbol{J} \in \mathcal{J}(X)^n$, judgment $J'_i \in \mathcal{J}(X)$, and $k \in \mathbb{R}$:

$$\text{if } k \in J'_i(x) \backslash J_i(x) \text{ and } J'_i(y) = J_i(y) \text{ for all } y \in X \text{ such that } y \neq x, \text{ then}$$
$$k \in F(\boldsymbol{J})(x) \Rightarrow k \in F(\boldsymbol{J}_{-i}, J'_i)(x)$$

The idea of monotonic support is that additional support should not harm what was already accepted. The mean rule is a clear example of a rule that does not have monotonic support, because the mean value of a collection does not depend on the support it gets.

Definition 3.9 (Directionality)**.** A rule F is *(weakly) directional* if, when one agent strictly changes their judgment, the outcome does not change in the other direction. For any variable $x \in X$, agent $i \in N$ and profile $\boldsymbol{J} \in \mathcal{J}(X)^n$ with $J_i(x) = [a_i, b_i]$ and $F(\boldsymbol{J})(x) = [\hat{a}, \hat{b}]$, the following hold:

(i) Consider $\boldsymbol{J'} = (\boldsymbol{J}_{-i}, J'_i)$ with $J'_i(x) = [a'_i, b_i]$ and $F(\boldsymbol{J'})(x) = [\hat{a}', \hat{b}']$. We have:

$$a'_i < a_i \Rightarrow \hat{a}' \leq \hat{a}$$
$$a'_i > a_i \Rightarrow \hat{a}' \geq \hat{a}$$

(ii) Consider $\boldsymbol{J'} = (\boldsymbol{J}_{-i}, J'_i)$ with $J'_i(x) = [a_i, b'_i]$ and $F(\boldsymbol{J'})(x) = [\hat{a}', \hat{b}']$. We have:

$$b'_i < b_i \Rightarrow \hat{b}' \leq \hat{b}$$
$$b'_i > b_i \Rightarrow \hat{b}' \geq \hat{b}$$

A rule F is *strongly directional* if, when one agent strictly changes their judgment, the outcome changes in the same direction. For any variable $x \in X$, agent $i \in N$ and profile $\boldsymbol{J} \in \mathcal{J}(X)^n$ with $J_i(x) = [a_i, b_i]$ and $F(\boldsymbol{J})(x) = [\hat{a}, \hat{b}]$, the following hold:

(i) Consider $\boldsymbol{J'} = (\boldsymbol{J}_{-i}, J'_i)$ with $J'_i(x) = [a'_i, b_i]$ and $F(\boldsymbol{J'})(x) = [\hat{a}', \hat{b}']$. We have:

$$a'_i < a_i \Rightarrow \hat{a}' < \hat{a}$$
$$a'_i > a_i \Rightarrow \hat{a}' > \hat{a}$$

(ii) Consider $\boldsymbol{J'} = (\boldsymbol{J}_{-i}, J'_i)$ with $J'_i(x) = [a_i, b'_i]$ and $F(\boldsymbol{J'})(x) = [\hat{a}', \hat{b}']$. We have:

$$b'_i < b_i \Rightarrow \hat{b}' < \hat{b}$$
$$b'_i > b_i \Rightarrow \hat{b}' > \hat{b}$$

Directionality requires that the outcome be in the form of a closed interval, but only in the preconditions.[1] So, although support-based rules (see Subsection 2.2.3) do not necessarily return an interval, they can still possess this property based on the profiles for which they do. For the unanimity rule, it can even happen that the change in bounds causes the maximum infimum to become larger than the minimum infimum, in which case the collective judgment value ends up empty.

Directionality is closely related to the notion of fairness because every agent has the power to influence the outcome by reporting a different judgment. On the one hand, every agent's vote counts. On the other hand, agents have the power to manipulate the outcome in their favor. The mean rule is an example of a rule that is obviously directional and thus manipulable.

[1] When extending the framework for, e.g., (half)-open intervals, or arbitrary sets, we could consider the respective infimums and supremums rather than the endpoints of the closed intervals.

Definition 3.10 (Enclosure)**.** F is *enclosing* if the collective judgment lies somewhere in the convex hull of all the reported values. For any profile $\boldsymbol{J} \in \mathcal{J}(X)^n$:

$$F(\boldsymbol{J})(x) \subseteq \text{hull}\left(\bigcup_{i \in N} J_i(x)\right)$$

If $\boldsymbol{J}$ and $F(\boldsymbol{J})$ are point-valued, this boils down to:

$$\min_{i \in N} J_i(x) \leq F(\boldsymbol{J})(x) \leq \max_{i \in N} J_i(x)$$

Enclosure characterizes an aggregation rule's trust in the agents, so to speak. If they claim that the value cannot be outside their reported range, then the rule should respect that. The axiom is linked to the notion of accuracy; groups of people tend to 'bracket the truth' (Larrick & Soll, 2006), meaning that the correct value, whatever that signifies in a certain context, often lies somewhere in the middle of the individual judgment values.

3.2 Axiomatic Analysis

This section analyzes how the aggregation rules and repair operations defined in Chapter 2 relate to the axioms defined in the previous section. We do this systematically by examining which axiomatic properties they satisfy or preserve. This allows us to assess their normative appeal and identify which rules are most appropriate under specific constraints or applications.

The results of our analysis are summarized in Table 3.1; it shows which properties are satisfied by each rule. We first discuss some observations from the table and follow up with more formal characterization results for the family of support-quota rules via the concept of winning coalitions. We then discuss the effects of applying repair operations, such as the convex hull and point repair, and examine whether the axiomatic properties of a rule are preserved under such transformations.

3.2.1 Analysis Results

Rule	Ano	Neu	Ind	Aut	Una	Sup	Ref	Mon	Dir	Enc
Mean	Y	G	G	N	L	Y	N	N	S	Y
Weighted Mean	N	G	G	N	L	Y	N	N	W	Y
Median	Y	G	G	N	L	Y	Y	N	W	Y
Min & Max	Y	G	G	N	L	Y	Y	Y	W	Y
Bound-Quota	Y	N	G	N	N	N	Y	N	W	Y
BQ uni	Y	G	G	N	N	N	Y	N	W	Y
BQ maj	Y	G	G	N	L	Y	Y	N	W	Y
Nomination	Y	L	L	Y	L	Y	Y	Y	W	Y
Unanimity	Y	L	L	Y	L	Y	Y	Y	W	Y
Mode	Y	G	G	Y	L	Y	Y	Y	W	Y
Support-Quota	Y	N	L	Y	N	N	Y	Y	W	Y
SQ uni	Y	L	L	Y	N	N	Y	Y	W	Y
SQ maj	Y	L	L	Y	L	Y	Y	Y	W	Y

Table 3.1: Axiomatic analysis results. The rows correspond to the judgment aggregation rules discussed in Section 2.2. The columns denote the axiomatic properties defined in Section 3.1. Legend: Y/N = yes/no, L/G/N = local/global/neither, S/W = strong/weak.

Firstly, we notice that all the rules satisfy the enclosure axiom. This is because we have only defined rules that are, for some reason, interesting to consider. Enclosure is one of the most self-evident axioms with respect to rationality. Any rule that returns values that fall out of the range of values reported by any of the agents seems like an odd and impractical rule. Similarly, we have not considered rules that are not weakly directional.

The simple mean rule (2.1) seems to be the only rule that satisfies strong directionality. In general, weighted mean rules (2.2) also have this property, but the axiom is violated when there is an agent whose judgment is assigned a weight of zero. For any weight function that assigns nonzero weights to all agents, the mean rules are strongly directional because every agent has the power, however small, to manipulate the outcome. The weight function is also the reason that, unlike all the other rules, weighted mean rules fail anonymity, unless the weights are uniform.

Unanimous refutation is the property that distinguishes ordinal rules from calculation-based rules, at least the ones presented here. Ordinal, as well as support-based, rules generally will not include values in the collective outcome that did not already appear in the input profile. In contrast, calculation-based rules return values that are the result of some calculation, regardless of whether those values were reported by an agent.

Unanimous support and local unanimity are satisfied by all rules, except for the quota rules (2.6 and 2.10). This is because the quota can extend the total number of agents, prohibiting values from being included in the collective judgment, even if they are unanimously reported. For lower quotas, the rules do satisfy these axioms.

The quota rules are also the only rules that are not neutral. This is because nonuniform quota rules are defined by quotas, which might not be the same for each variable. We have only introduced the notion of uniformity for quota rules, but any aggregation rule that determines the collective judgment value differently for each variable is not neutral.

We can see by their definitions that the uniform rules introduced are both globally neutral and globally independent. Most support-based rules additionally satisfy local neutrality and local independence. The family of support-quota rules is even locally independent by purpose.

Proposition 3.1. *Support-quota rules are locally independent.*

Proof. Consider a support-quota rule F_{sq} defined by a quota function q. For any variable $x \in X$, profiles $\boldsymbol{J}, \boldsymbol{J'} \in \mathcal{J}(X)^n$ and $k \in \mathbb{R}$, we have $k \in F_{\text{sq}}(\boldsymbol{J})(x) \Leftrightarrow \# N^{\boldsymbol{J}}_{x,k} \geq q(x)$ and, similarly, $k \in F_{\text{sq}}(\boldsymbol{J'})(x) \Leftrightarrow \# N^{\boldsymbol{J'}}_{x,k} \geq q(x)$. Since $k \in J_i(x) \Leftrightarrow k \in J'_i(x)$ for all $i \in N$ means that $N^{\boldsymbol{J'}}_{x,k} = N^{\boldsymbol{J}}_{x,k}$, it implies that $k \in F_{\text{sq}}(\boldsymbol{J})(x) \Leftrightarrow k \in F_{\text{sq}}(\boldsymbol{J'})(x)$. Thus, F_{sq} is locally independent. $\square$

A similar proof shows that uniform support-quota rules are locally neutral.

The mode rule (2.9) is a support-based rule that does not satisfy local neutrality and independence, something it has in common with calculation-based and ordinal rules. None of these rules can determine the collective judgment values locally, because they rely on other reported numbers; calculation-based rules need all reported numbers respective to the calculation, ordinal rules need all reported numbers to order the judgments, and the mode rule needs all reported numbers to determine which number receives the most support. This is made explicit in the example below.

Example 3.1. Consider a set of 10 agents such that $N = A \uplus B \uplus C$ with $|A| = 4$ and $|B| = |C| = 3$. Let $\boldsymbol{J}$ be the following point-valued profile on $X = \{x, y\}$, for $i \in N$:

$$J_i(x) = \begin{cases} \mathbf{2} & \text{if } i \in A \\ \mathbf{1} & \text{if } i \in B \\ \mathbf{3} & \text{if } i \in C \end{cases} \qquad\qquad J_i(y) = \begin{cases} \mathbf{2} & \text{if } i \in A \\ \mathbf{3} & \text{if } i \in B \cup C \end{cases}$$

Let us evaluate the collective judgment values for a calculation-based rule, an ordinal rule, and the mode rule:

	x	y
$F_{\text{mean}}(\boldsymbol{J})$	**2**	**2.6**
$F_{\text{min}}(\boldsymbol{J})$	**1**	**2**
$F_{\text{mode}}(\boldsymbol{J})$	**2**	**3**

All these rules fail the local neutrality axiom for $k = 2$; although $N^{\boldsymbol{J}}_{x,2} = N^{\boldsymbol{J}}_{y,2} = A$, the number 2 is not collectively accepted for either both or neither of the variables. If we consider a profile $\boldsymbol{J'}$ such that $J'_i(x) = J_i(y)$ for all $i \in N$, it is obvious that the rules are not locally independent either.

Even though the autonomy axiom and the local neutrality and independence axioms look similar, the mode rule proves that they represent different properties. The mode rule does not satisfy the latter, because it has to compare the amount of support to determine which number receives the most support. Although other support-based rules also depend on the notion of support, they do not care about how it compares to that of other numbers. In contrast, unlike calculation-based and ordinal rules, the mode rule is autonomous, as are all other support-based rules. In fact, it seems to be a necessary property for support-based rules. This is because the notion of support is not defined by the objects that are being supported. That is, the numbers could have been colors, types of fruit, or names of people, for that matter. The example below shows how autonomy is, or is not, satisfied by some rules.

Example 3.2. Consider a profile $\boldsymbol{J} = (J_1, J_2, J_3)$ with the following judgments on $x \in X$:

$$J_1(x) = [0, 4]$$
$$J_2(x) = [3, 10]$$
$$J_3(x) = [6, 8]$$

These judgments induce the following collective judgment values for x:

$$F_{\cup}(\boldsymbol{J})(x) = [0, 10] \qquad F_{\text{mean}}(\boldsymbol{J})(x) = \left[3, 7\tfrac{1}{3}\right]$$
$$F_{\cap}(\boldsymbol{J})(x) = \varnothing \qquad F_{\text{med}}(\boldsymbol{J})(x) = [3, 8]$$
$$F_{\text{mode}}(\boldsymbol{J})(x) = [3, 4] \cup [6, 8] \qquad F_{\text{max}}(\boldsymbol{J})(x) = [6, 10]$$
$$F_{\text{sq,maj}}(\boldsymbol{J})(x) = [3, 4] \cup [6, 8] \qquad F_{\text{bq,maj}}(\boldsymbol{J})(x) = [3, 8]$$

The support-based rules on the left all satisfy the autonomy axiom. The rules on the right are calculation-based and ordinal rules that fail the axiom for $k = 5$ and $k' = 9$; although $N^{\boldsymbol{J}}_{x,5} = N^{\boldsymbol{J}}_{x,9} = \{2\}$, the collective judgment values for x do not include both or neither of the numbers.

As the name already suggests, monotonic support is another property that can be satisfied by support-based rules. Non-monotonic support-based rules do exist, but we have not discussed them. For example, the rule that returns the values that receive the least amount of support, i.e., the converse of the mode rule, is support-based but not monotonically supporting. The opposite, a rule that is monotonically supporting but not support-based, is rarer. If a value receives additional support, there must have been an agent that changed their judgment, which influences the calculation or order of the judgment values. So, calculation-based and ordinal rules might discard a value that was previously accepted. The min and max rules are two exceptions. However, their satisfying monotonic support is more a side effect of the rules' mechanics than a property that is intuitively ascribed to them. It is left to the reader to check this.

3.2.2 Characterization Results

In JA, the concept of *winning coalitions* is a means to provide the axiomatic characterization of aggregation rules (Endriss, 2016). We adapt the concept to our numerical framework by interpreting each pair $(x, k) \in X \times \mathbb{R}$ as a statement: "k is a possible value for variable x." A winning coalition for (x, k) is then a set of agents with the collective power to have k be included in the collective judgment for x.

Definition 3.11 (Winning Coalition)**.** Let F be an aggregation rule. For every variable $x \in X$ and value $k \in \mathbb{R}$, we call $\mathcal{W}_{x,k} \subseteq N$ the family of *winning coalitions* for (x, k) such that for any profile $\boldsymbol{J} \in \mathcal{J}(X)^n$, we have:

$$k \in F(\boldsymbol{J})(x) \text{ if and only if } N^{\boldsymbol{J}}_{x,k} \in \mathcal{W}_{x,k}$$

The notion of winning coalitions is related to the property of local independence (3.6). In fact, locally independent rules can be alternatively defined by them:

Lemma 3.2 (Local Independence)**.** *An aggregation rule F is locally independent if and only if there exists a family of winning coalitions $\mathcal{W}_{x,k} \subseteq 2^N$, one for every $x \in X$ and $k \in \mathbb{R}$, such that for all profiles $\boldsymbol{J} \in \mathcal{J}(X)^n$, it is the case that $k \in F(\boldsymbol{J})(x)$ if and only if $N^{\boldsymbol{J}}_{x,k} \in \mathcal{W}_{x,k}$.*

For locally independent rules, we can reformulate other axioms in terms of their associated winning coalitions.

Lemma 3.3 (Winning Coalitions)**.** *Let F be a locally independent aggregation rule, and, for every $x \in X$ and $k \in \mathbb{R}$, let $\mathcal{W}_{x,k} \subseteq 2^N$ be the corresponding family of winning coalitions. Then:*

(i) F is anonymous if and only if $\mathcal{W}_{x,k}$ is closed under equinumerosity; $C \in \mathcal{W}_{x,k}$ and $|C| = |C'|$ imply $C' \in \mathcal{W}_{x,k}$ for all $C, C' \subseteq N$ and all $x \in X$ and $k \in \mathbb{R}$.

(ii) F is a dictatorship if and only if there exists an $i \in N$ such that for all $x \in X$ and $k \in \mathbb{R}$, $C \in \mathcal{W}_{x,k} \Leftrightarrow i \in C$.

(iii) F is unanimously supporting if and only if $N \in \mathcal{W}_{x,k}$ for all $x \in X$ and $k \in \mathbb{R}$.

(iv) F is unanimously refuting if and only if $\varnothing \notin \mathcal{W}_{x,k}$ for all $x \in X$ and $k \in \mathbb{R}$.

(v) F is locally neutral if and only if $\mathcal{W}_{x,k} = \mathcal{W}_{y,k}$ for all $x, y \in X$ and $k \in \mathbb{R}$.

(vi) F is autonomous if and only if $\mathcal{W}_{x,k} = \mathcal{W}_{x,k'}$ for all $x \in X$ and $k, k' \in \mathbb{R}$.

(vii) F has monotonic support if and only if $\mathcal{W}_{x,k}$ is upward-closed; $C \in \mathcal{W}_{x,k}$ and $C \subseteq C'$ imply $C' \in \mathcal{W}_{x,k}$ for all $C, C' \subseteq N$, $x \in X$ and $k \in \mathbb{R}$.

Proof. The claims follow from the relevant definitions. See Appendix B for the full proof. □

Winning coalitions, in the way they are defined here, are only relevant for locally independent rules because they do not capture the numerical calculations or interval properties. However, at least for these rules, they are a very powerful characterization tool. Next, we present some characterization results on the family of support-quota rules.

Theorem 3.4. *An aggregation rule is anonymous, locally independent, autonomous, and has monotonic support if and only if it is a support-quota rule (2.10).*

Proof. The right-to-left direction follows immediately from checking that every support-quota rule clearly has those properties. The left-to-right direction follows from Lemma 3.3: By local independence and autonomy, we can decide on each number for each variable. That is, we consider the support $N^{\boldsymbol{J}}_{x,k}$ for each x and k independently. By monotonic support and anonymity, the family of winning coalitions is upward-closed and closed under equinumerosity. That is, we are interested in the amount of support $\#N^{\boldsymbol{J}}_{x,k}$ and any greater amount. Taken together, this means that for every variable $x \in X$, there exists a quota $q(x)$ such that $k \in F(\boldsymbol{J})(x)$ if and only if $\#N^{\boldsymbol{J}}_{x,k} \geq q(x)$. □

This result uniquely identifies the properties that are both necessary and sufficient for a rule to be a support-quota rule. Considering additional properties yields the following corollaries as immediate consequences of Theorem 3.4.

Corollary 3.5. *An aggregation rule is anonymous, locally neutral, locally independent, autonomous, and has monotonic support if and only if it is a uniform support-quota rule.*

A support-quota rule is locally neutral if and only if it is a uniform support-quota rule, because treating all variables equally means the quotas must be the same for all variables.

Corollary 3.6. *An aggregation rule is anonymous, locally independent, autonomous, unanimously refuting, and has monotonic support if and only if it is a support-quota rule with quotas of at least 1.*

Corollary 3.7. *An aggregation rule is anonymous, locally independent, autonomous, unanimously supporting, and has monotonic support if and only if it is a support-quota rule with quotas of at most n.*

With quotas between 1 and n, the support-quota rule will be unanimously supporting and refuting. Special cases include the nomination rule for $\lambda = 1$, the unanimity rule for $\lambda = n$ and the majority support rule for $\lambda = \lceil \frac{n+1}{2} \rceil$.

There are many interesting rules other than locally independent rules to apply when aggregating numerical judgments. Unfortunately, we have not yet found a tool that makes the axiomatic characterization of other rules as accessible and intuitive as winning coalitions.

3.2.3 Preservation of Properties

We have analyzed what properties certain aggregation rules have. However, we have not discussed how the application of repair operations impacts the properties of the composite procedure. Because operations modify the outputs of aggregation rules to meet practical requirements, such as returning intervals or single points, they can affect the axiomatic properties as well.

Simply put, the properties of the whole are the common properties of its parts. Of course, we have not defined properties for repair operations, but we can interpret the composite properties intuitively. For example, it does not make sense to call a repair operation anonymous, because repair operations do not handle information about agents. What we can say is that the nomination rule modified by the convex core operation is anonymous because any permutation of the agents in the input profile does not result in a different outcome. That is, taking the convex core of an outcome does not make the compound procedure less anonymous than the original aggregation rule. In fact, precisely because repair operations do not handle information about agents, all repair operations preserve anonymity.

Point repair operations revoke the property of unanimous support and local unanimity because they will discard all but one value, even if multiple were unanimously supported. Similarly, the convex hull operation revokes the property of unanimous refutation, since it includes values not originally reported, in order to close the gaps. In contrast, the convex core operation is more conservative, so to speak, and preserves unanimous refutation, as it only includes values that were already present in the outcome.

Thus, while repair operations improve the practical feasibility of aggregation outputs, they can reduce the normative appeal of aggregation procedures by weakening or destroying some desirable theoretical properties. This trade-off between fairness and accuracy versus feasibility must be evaluated relative to the application context.

Chapter 4

Integrity Constraints

Consider the following scenario. Alex and Bailey are supervisors for a group of scouts. They are going on a hike, and they need to make sure they have enough food for everyone. They are planning to bring x bananas, y boxes of raisins, and z bags of nuts. Independently, they make an estimate of how many snacks to bring. To check whether they expect the kids to prefer sweet snacks over salty snacks, we can check whether their judgments satisfy the constraint $x + y > z$. As we can see from the profile in the figure below, J_A is consistent with the constraint, and J_B is not. This means that Alex thinks the kids prefer sweet snacks, and Bailey thinks they prefer salty snacks.

	x	y	z
J_A	3	5	4
J_B	4	1	7

Figure 4.1: Example profile $\boldsymbol{J} = (J_A, J_B)$ on agenda $X = \{x, y, z\}$.

The example scenario shows how a constraint can be used to extract useful information from a judgment profile. Constraints are broadly applicable. They can be used to analyze the connectedness of the variables, as in the example above; to externally impose a structure on the agenda, for example, to emulate an algebraic agenda; or to restrict the input profile, for example, by allowing only bivalent judgments.

In general, integrity constraints are used to describe what makes a judgment rational or feasible. For example, in the context of classical logic, only bivalent judgments can be considered rational; reporting the value 17 for a propositional variable makes no sense. Similarly, for a problem of preference aggregation, we could impose a constraint that encodes the antisymmetry or transitivity of a preference ordering. Standard judgment aggregation involves an agenda of propositional formulas. However, it is possible to reduce the language of the agenda, e.g., only allowing less complex variables, but preserve the semantics by imposing integrity constraints on the relations between those simpler variables. Grandi and Endriss (2010) define a binary aggregation framework that incorporates integrity constraints in the language of propositional logic. Because NJA allows for judgments to be more nuanced than binary judgments, we can consider more complex and meaningful constraints than those expressed as propositional formulas.

4.1 Rationality and Feasibility Constraints

Endriss (2018) distinguishes between rationality and feasibility constraints. The former restrict the input profile. They represent the supposed rationality of the individual agents. The latter restrict the output. They represent what we consider to be a feasible collective decision on the values of the issues under consideration. An integrity constraint is not by definition a rationality constraint or a feasibility constraint. Rather, the function of a constraint is based on the application. The set of feasibility constraints could contain the same constraints as the set of rationality constraints, one could be more demanding than the other, or they could contain entirely different constraints.

To ensure that the input profile is, for example, point-valued, we would impose the following rationality constraint on each agent's judgment J:

$$\forall x \in X : \exists k \in \mathbb{R} : J(x) = \{k\} \tag{4.1}$$

If, under some aggregation rule, the outcome is point-valued as well, we say the outcome is collectively rational. Collective rationality is used to describe when rationality constraints imposed on individuals judgment are preserved in the collective judgment. We can also imagine a scenario where agents are not limited to reporting point-valued judgments, but the application does call for a point-valued decision. In that case, we would impose the point-valuedness constraint on the collective judgment as a feasibility constraint.

We can formalize any of the judgment function properties defined in Definition 2.1 as an integrity constraint. However, it is more interesting to discuss constraints that tell us something about how variables or their judgment values are connected. For example, in the context of sustainability, we might want to represent emission standards, fishing quotas, or exposure limits. We can formulate a constraint that says to cut total CO_2-emissions by at least p kg:

$$\sum_{x \in X} J(x) \geq p \tag{4.2}$$

In this constraint, each variable $x \in X$ represents a measure to reduce CO_2 such that $J(x)$ denotes the amount that measure x supposedly contributes to the total weight of CO_2.

Similarly, we can formulate a constraint to impose an upper bound on the total of the values:

$$\sum_{x \in X} J(x) \leq p \tag{4.3}$$

A possible reading of this constraint is that at most p acres of forest can be cut down, where each variable $x \in X$ represents the surface area of a woodland region.

These constraints are not limited to the sustainability context. The parameter p can represent any bound on the total of the values, and X can represent any set of resources, products, or activities. In the context of budgeting, p could be the budget and $\sum_{x \in X} J(x)$ the total cost of the projects in X. Or, for human resource management, p could be the hours available to all employees and $\sum_{x \in X} J(x)$ the total hours needed for tasks in X.

4.2 Consistency and Safety

Constraints can tell us something about how variables are connected or not. This is not determined by the contents of the judgments, but rather the contents of the agenda. In standard judgment aggregation, the logical consistency of a judgment is determined by the set of formulas that are judged true. This notion of consistency is thus tied to the propositional nature of the agenda-items. Similarly, the notion of probabilistic consistency in opinion pooling is tied to the probabilistic structure of the agenda. Since the NJA framework is constructed for general agendas, we can no longer rely on their structure to inform us about the consistency of judgments. However, using integrity constraints, we can describe consistency on a higher level. Informally, we say that a judgment is consistent if it complies with the integrity constraints imposed on it. To make this formal, we first need to define what it means for a judgment to comply with a constraint.

Definition 4.1 (Satisfaction)**.** A judgment J *satisfies* a constraint Γ, denoted $J \vDash \Gamma$, if Γ is true when the variables in Γ are assigned the values given by J.

We can call a judgment function $J : X \to 2^{\mathbb{R}}$ an assignment because it maps variables to sets of values. If we were to use NJA to model a judgment aggregation problem based on propositional logic, the agenda-items would represent propositional formulas, and a judgment function would represent an assignment from propositional variables to truth values, i.e., a valuation. Similarly, for a probabilistic aggregation problem, a judgment function would represent an assignment from probabilistic events to probabilities, i.e., a probability measure.

Definition 4.2 (Consistency)**.** A judgment J is *consistent* with a set of constraints Δ if J satisfies all the constraints in Δ, i.e., $J \vDash \Gamma$ for all $\Gamma \in \Delta$, and is otherwise *inconsistent*. We simply say that J is consistent if Δ is clear from context. We extend this definition to say that a profile $\boldsymbol{J}$ is consistent (with a set of constraints Δ) if every judgment J_i for $i \in N$ is consistent (with Δ).

Consider the following point-valued profile:

	x	y
J_1	**1**	**4**
J_2	**4**	**1**
J_3	**1**	**4**
$F_{\text{mean}}(\boldsymbol{J})$	**2**	**3**

Figure 4.2: Example profile.

It is easy to see that each individual judgment is consistent with the constraints $x+y = 5$ and $x \times y = 4$. However, the collective judgment is only consistent with the former. We say that the latter constraint is not *safe* for the mean rule. That is, the mean rule can yield an inconsistent outcome even when the profile is consistent.

Definition 4.3 (Safety)**.** A constraint Γ is *safe* for an aggregation rule F if for any profile $\boldsymbol{J} \in \mathcal{J}(X)^n$ that satisfies Γ, the collective outcome $F(\boldsymbol{J})$ also satisfies Γ, that is, $\boldsymbol{J} \vDash \Gamma$ implies $F(\boldsymbol{J}) \vDash \Gamma$.

4.3 Arithmetic Constraints

In this section, we will only consider constraints for point-valued judgments. Although most of the operations are already defined for set-valued judgments (see Appendix A), it has not been investigated whether the results can be extended for set-valued judgments in general.

In the previous section, we demonstrated via a counterexample (see Figure 4.2) that the multiplicative constraint is not safe for the mean rule. The example indicates that the additive constraint is safe for the mean rule, but to show that, we need a proper proof. Firstly, we note that $x + y = 5$ is just an instance of an additive constraint. We can construct more general arithmetic constraints on variables. For some k designated variables $x_1, \ldots, x_k \in X$ and constants $\alpha_1, \ldots, \alpha_k, \beta \in \mathbb{R}$, we have:

$$\alpha_1 x_1 + \cdots + \alpha_k x_k = \beta \qquad \text{(addition)}$$

$$x_1 \times \cdots \times x_k = \beta \qquad \text{(multiplication)}$$

Note that in the multiplicative constraint, all $\alpha_1, \ldots, \alpha_k$ are subsumed by β.

We can show that even this more general addition constraint is safe for the mean rule for any $k \leq |X|$.

Proposition 4.1. *For any variables $x_1, \ldots, x_k \in X$ and constants $\alpha_1, \ldots, \alpha_k, \beta \in \mathbb{R}$, any additive constraint of the form $\alpha_1 x_1 + \cdots + \alpha_k x_k = \beta$ is safe for the mean rule.*

Proof. Given a consistent profile $\boldsymbol{J} = (J_1, \ldots, J_n)$, the mean rule yields an outcome with $F_{\text{mean}}(\boldsymbol{J})(x) = \frac{\sum_{i=1}^n J_i(x)}{n}$ for each $x \in X$. Since each individual agent is consistent, we have $\alpha_1 J_i(x_1) + \cdots + \alpha_k J_i(x_k) = \beta$ for all $i \in N$. Therefore, we get

$$\begin{aligned}
\alpha_1 F_{\text{mean}}(\boldsymbol{J})(x_1) + \cdots + \alpha_k F_{\text{mean}}(\boldsymbol{J})(x_k) &= \alpha_1 \frac{\sum_{i=1}^n J_i(x_1)}{n} + \cdots + \alpha_k \frac{\sum_{i=1}^n J_i(x_k)}{n} \\
&= \frac{\sum_{i=1}^n \alpha_1 J_i(x_1)}{n} + \cdots + \frac{\sum_{i=1}^n \alpha_k J_i(x_k)}{n} \\
&= \frac{\sum_{i=1}^n (\alpha_1 J_i(x_1) + \cdots + \alpha_k J_i(x_k))}{n} \\
&= \frac{\sum_{i=1}^n \beta}{n} \\
&= \beta
\end{aligned}$$

□

This result also makes sense intuitively. The parameters of an affine transformation[1] are constants and will be averaged throughout the agents for each variable and simply returned again by the mean. This is a property specific to the mean rule. The median rule, for example, does not have this property.

Example 4.1. Additive constraints are not safe for the median rule in general. Consider the constraint $x + y - z = 0$. We construct the following point-valued profile:

	x	y	z
J_1	**2**	**1**	**3**
J_2	**2**	**3**	**5**
J_3	**1**	**3**	**4**
$F_{\text{med}}(\boldsymbol{J})$	**2**	**3**	**4**

Clearly, each individual judgment is consistent, but the collective judgment is not.

Similarly, we can give a counterexample that shows that the general multiplication constraint is not safe for the median rule.

Example 4.2. Multiplicative constraints are not safe for the median rule in general. Consider the constraint $x \times y \times z = 1$. We construct the following profile:

	x	y	z
J_1	**1**	$\mathbf{\frac{1}{2}}$	**2**
J_2	**1**	**3**	$\mathbf{\frac{1}{3}}$
J_3	**3**	**2**	$\mathbf{\frac{1}{6}}$
$F_{\text{med}}(\boldsymbol{J})$	**1**	**2**	$\mathbf{\frac{1}{3}}$

Clearly, each individual judgment is consistent, but the collective judgment is not.

If we rewrite $x+y-z=0$ as $x+y=z$, it turns out that Example 4.1 is a generalization of the disjunctive version of the doctrinal paradox (Figure 1.2). For the standard doctrinal paradox, that is, the conjunctive version, we can consider the constraint $x \times y = z$ as its generalization.

[1] A composition of a linear transformation and a translation (see Appendix A).

	x	y	$x \times y$
J_1	**1**	**0**	**0**
J_2	**0**	**1**	**0**
J_3	**1**	**1**	**1**
$F_{\mathrm{med}}(\boldsymbol{J})$	**1**	**1**	**0**

(a) Multiplication (conjunctive paradox in JA)

	x	y	$x + y$
J_1	**1**	**0**	**1**
J_2	**0**	**1**	**1**
J_3	**0**	**0**	**0**
$F_{\mathrm{med}}(\boldsymbol{J})$	**0**	**0**	**1**

(b) Addition (disjunctive paradox in JA)

Figure 4.3: Doctrinal paradoxes in NJA.

In the figure above, we consider an agenda with three variables x, y, and z such that z represents either $x \times y$ or $x + y$. Thus far, we have only seen examples of how the median rule does not respect arithmetic constraints with three variables. Clearly, any arithmetic constraint involving even more variables will also not be safe for the median rule. However, for smaller values of k, it turns out that the addition and multiplication constraints are safe for the median rule, albeit under certain conditions. For $k = 1$, it is easy to see that the safety of the constraints follows immediately from the definition of the median rule. The median rule returns the median value of the variable, which was reported by some agent $i \in N$. If their judgment is consistent, then clearly the collective judgment consistent as well. For $k = 2$, the multiplication constraint is only safe for the median rule when β is nonnegative and the number of agents is odd. If the profile in Figure 4.2 consisted only of J_1 and J_2, for example, the collective judgment using the median rule would not be consistent with $x \times y = 4$. This is because the median rule uses the upper median (see Definition 2.2).

Proposition 4.2. *For an odd number of agents, any multiplicative constraint of the form $x \times y = \beta$ with $\beta \geq 0$ is safe for the median rule.*

Proof. Let $\boldsymbol{J} \in \mathcal{J}(X)^n$ be a consistent point-valued profile with odd n. Given that all individual judgments satisfy the constraint, it must be that for every $i, j \in N$ such that $J_i(x) < J_j(x)$, we have $J_i(y) > J_j(y)$. If we order the agents in ascending order for their reported value on x, then they must be in descending order for their reported value on y. Because n is odd, there is an agent $i \in N$ such that J_i is the middle judgment in the ordering. By definition of the median rule, the collective judgment values for x and y are $J_i(x)$ and $J_i(y)$, respectively, so the collective judgment is consistent. Thus, the constraint is safe for the median rule. □

With nearly the same proof, we can show that, for an odd number of agents, any additive constraint of the form $\alpha_1 x + \alpha_2 y = \beta$ is also safe for the median rule, even if β is negative. It is left to the reader to verify this. To show that the multiplicative constraint is not safe for negative β, we provide the following counterexample.

Example 4.3. A multiplication constraint of the form $x \times y = \beta$ with $\beta < 0$ is not safe for the median rule. Consider the constraint $x \times y = -12$. We construct the following profile:

	x	y
J_1	**3**	**−4**
J_2	**−2**	**6**
J_3	**4**	**−3**
$F_{\mathrm{med}}(\boldsymbol{J})$	**3**	**−3**

Each individual judgment is consistent, but the collective judgment is not.

If we impose additional constraints, we can work around the paradox. For example, if we require $x \geq 0$, the conjoint constraints are safe for the median rule if β is negative, even for an even number of agents.

Proposition 4.3. *Any multiplicative constraint of the form $x \times y = \beta$ with $\beta < 0$, together with the constraint $x \geq 0$, is safe for the median rule.*

Proof. Let $\boldsymbol{J} \in \mathcal{J}(X)^n$ be a consistent point-valued profile. We order the agents in ascending order for their reported value on x. Given that all individual judgments satisfy both constraints and β is negative, it must be that all judgment values on y are negative. In fact, for every $i, j \in N$ such that $J_i(x) < J_j(x)$, we also have $J_i(y) < J_j(y)$, i.e., the agents are also in ascending order for their reported value on y. By definition of the median rule, the collective judgment values for x and y are the respective judgment values reported by the (upper) median agent in the ordering, so the collective judgment is consistent. Thus, the conjoint constraints are safe for the median rule. □

The constraint $x \geq 0$ is an inequality constraint, as opposed to an equality constraint. Another inequality constraint for some $x, y \in X$ is the following:

$$x > y \qquad \text{(pairwise inequality)}$$

We could generalize this constraint to include more variables or constants, just as we did for the addition and multiplication constraints, but the results are similar.

Ordering the profile is a convenient and intuitive proof strategy. To show that pairwise inequality is safe for the median rule, we apply this method again. However, the proof is a little more complex than the others we have seen so far.

Proposition 4.4. *Any constraint of the form $x > y$ is safe for the median rule.*

Proof. Let $\boldsymbol{J} \in \mathcal{J}(X)^n$ be a consistent point-valued profile. We order the agents in ascending order for their reported value on x. By definition of the median rule, $F_{\text{med}}(\boldsymbol{J})(x) = J_i(x)$ such that $i \in N$ is the middle agent in the profile ordering if n is odd, or right below the middle otherwise. Because J_i is consistent, we have $J_i(x) > J_i(y)$. There is also some $j \in N$ such that $F_{\text{med}}(\boldsymbol{J})(y) = J_j(y)$ and $J_j(x) > J_j(y)$.

If $i = j$ or if agent j is above agent i in the ordering, we have $J_i(x) \geq J_j(x)$, which implies that $J_i(x) > J_j(y)$, hence $F_{\text{med}}(\boldsymbol{J})(x) > F_{\text{med}}(\boldsymbol{J})(y)$. If agent j is below agent i in the ordering, strictly more than half of the agents are above agent j in the ordering. For $J_j(y)$ to be the median value for y, there must then be at least one agent k somewhere above agent i in the ordering, i.e., $J_k(x) \leq J_i(x)$, such that $J_k(y) \geq J_j(y)$. By consistency of J_k, we know that $J_k(x) > J_k(y)$, so $J_k(x) > J_j(y)$. It follows that $J_i(x) > J_j(y)$, so $F_{\text{med}}(\boldsymbol{J})(x) > F_{\text{med}}(\boldsymbol{J})(y)$.

In any case, we have $F_{\text{med}}(\boldsymbol{J})(x) > F_{\text{med}}(\boldsymbol{J})(y)$, so $F_{\text{med}}(\boldsymbol{J})$ is consistent. Thus, the constraint is safe for the median rule. □

In this section, we have only considered results for point-valued input and output. More research is needed to provide results on set-valued judgments. Fortunately, most of the operations in this section are already applicable to set-valued judgments. For example, addition and multiplication for sets are defined in Appendix A. However, the phrase $x > y$ only has meaning for point-valued judgments. Unless explicitly defined, the inequality of two sets is ambiguous.

Inequality of sets could describe a cardinality comparison; S has more elements than T. Or it could describe a lexicographic order between sets; the elements in S are greater than those in T. In yet another interpretation, it could denote the subset relation; S is a superset of T. Since point-valued judgments are special cases of set-valued judgments, we implicitly already interpret a lexicographic ordering, but regardless of the meaning of the inequality operator, it does not influence the correctness of the proofs for Propositions 4.2 and 4.4. The reason we have not adapted them for set-valued judgments is that the rules introduced in Section 2.2 are only defined for interval input; this includes point-valued input but not set-valued input.

4.4 Algebraic Semantics

Another use of integrity constraints is to impose certain formal semantics. Suppose we want to model an aggregation problem in the context of propositional logic. A propositional aggregation problem requires agents to judge the truth values of a set of propositions. Because the agenda in NJA has no internal structure, the logical semantics do not come for free. However, we can introduce integrity constraints that encode them. Given a propositional language $\mathcal{L}$, the agenda $X \subseteq \{x_\varphi : \varphi \in \mathcal{L}\}$ consists of variables that are indexed by the formulas in $\mathcal{L}$ they represent. Of course, this is purely syntactical; the variables are still meaningless in themselves. Let $\mathcal{T}$ be the set of possible truth values. We construct the following constraints for truth-functional operations on judgments:

$$\begin{aligned}
&\forall x_\varphi \in X: && x_{\neg\varphi} = \mathbf{1} - x_\varphi && \text{(Negation, subtraction)}\\
&\forall x_\varphi, x_\psi \in X: && x_{\varphi\wedge\psi} = \min(x_\varphi, x_\psi) && \text{(Conjunction, minimum)}\\
&\forall x_\varphi, x_\psi \in X: && x_{\varphi\vee\psi} = \max(x_\varphi, x_\psi) && \text{(Disjunction, maximum)}\\
& && x_\perp = \mathbf{0} && \text{(Falsum point)}\\
& && x_\top = \mathbf{1} && \text{(Verum point)}
\end{aligned}$$

Note that all the operations in the constraints are set operations (see Appendix A). We call this set of constraints $\Delta_{\mathcal{L},\min}$. A judgment is logically consistent if it satisfies all the constraints in $\Delta_{\mathcal{L},\min}$. If we set $\mathcal{T} = \{0, 1\}$, we can model aggregation problems involving classical logic. To model other logical systems, we might need to adjust the constraints. For example, if we set $\mathcal{T} = [0, 1]$ for the constraints above, we get a fuzzy logic, but if we want Gödel logic (1932), we need a different constraint for negation and an additional one for the implication operation.

Example 4.4. In the table below, each row represents an example judgment function that is logically consistent with $\Delta_{\mathcal{L},\min}$ for the set of truth values $\mathcal{T}$.

$\mathcal{T}$	x_p	x_q	$x_{\neg p}$	$x_{p\wedge q}$	$x_{p\vee q}$
$\{0,1\}$	$\mathbf{0}$	$\mathbf{1}$	$\mathbf{1}$	$\mathbf{0}$	$\mathbf{1}$
$\{0,\frac{1}{2},1\}$	$\{0,\frac{1}{2}\}$	$\{\frac{1}{2},1\}$	$\{\frac{1}{2},1\}$	$\{0,\frac{1}{2}\}$	$\{\frac{1}{2},1\}$
$[0,1]$	$\{0,\frac{3}{4}\}$	$\{\frac{1}{4},1\}$	$\{\frac{1}{4},1\}$	$\{0,\frac{1}{4},\frac{3}{4}\}$	$\{\frac{1}{4},\frac{3}{4},1\}$
$[0,1]$	$[0,\frac{2}{3}]$	$[\frac{1}{3},1]$	$[\frac{1}{3},1]$	$[0,\frac{2}{3}]$	$[\frac{1}{3},1]$

An alternative to using the pairwise maximum and minimum for conjunction and disjunction, respectively, is using the pairwise product and probabilistic sum. We use the pairwise probabilistic sum, and not the Minkowski sum, to normalize for values that exceed the maximum truth value of 1.[2] For the operations to work properly, we then also have to change the subtractive negation to a negation by set difference. Let $\Delta_{\mathcal{L},\times}$ denote the following set of constraints together with the falsum and verum constraints as in $\Delta_{\mathcal{L},\min}$:

$$\begin{aligned}
&\forall x_\varphi \in X: && x_{\neg\varphi} = \mathcal{T} \setminus x_\varphi && \text{(Negation, difference)}\\
&\forall x_\varphi, x_\psi \in X: && x_{\varphi\wedge\psi} = x_\varphi \times x_\psi && \text{(Conjunction, product)}\\
&\forall x_\varphi, x_\psi \in X: && x_{\varphi\vee\psi} = x_\varphi \mathbin{\hat{+}} x_\psi && \text{(Disjunction, prob. sum)}
\end{aligned}$$

[2]This is presuming that the truth values range between 0 and 1. For other sets of truth values, normalization might not be necessary.

Example 4.5. In the table below, each row represents an example judgment function that is logically consistent with $\Delta_{\mathcal{L},\times}$ for the set of truth values $\mathcal{T}$.

$\mathcal{T}$	x_p	x_q	$x_{\neg p}$	$x_{p\wedge q}$	$x_{p\vee q}$
$\{0,1\}$	**0**	**1**	**1**	**0**	**1**
$[0,1]$	$\{0,\frac{1}{2}\}$	$\{\frac{1}{2},1\}$	$(0,\frac{1}{2})\cup(\frac{1}{2},1]$	$\{0,\frac{1}{4},\frac{1}{2}\}$	$\{\frac{1}{2},\frac{3}{4},1\}$
$[0,1]$	$[0,\frac{2}{3}]$	$[\frac{1}{3},1]$	$(\frac{2}{3},1]$	$[0,\frac{2}{3}]$	$[\frac{1}{3},1]$

Note that classical logic can be modeled with both $\Delta_{\mathcal{L}}$ and $\Delta_{\mathcal{L},\times}$, as long as $\mathcal{T}=\{0,1\}$.

If we set $\mathcal{T}=[0,1]$, we can interpret $\Delta_{\mathcal{L},\times}$ to encode probabilistic semantics. We construct the following constraints for point-valued judgments to encode the Kolmogorov axioms (1950) with respect to $\mathcal{L}$:

$$\forall x_\varphi \in X: \qquad x_\varphi \geq 0 \qquad \text{(Non-negativity)}$$
$$\forall x_\varphi \in X \text{ s.t. } \vDash \varphi: \qquad x_\varphi = 1 \qquad \text{(Tautologies)}$$
$$\forall x_\varphi, x_\psi \in X \text{ s.t. } \vDash \neg(\varphi\wedge\psi): \qquad x_{\varphi\vee\psi} = x_\varphi + x_\psi \qquad \text{(Finite additivity)}$$

A judgment is probabilistically consistent if it satisfies all constraints in $\Delta_{\mathcal{L},\times}$ and the Kolmogorov constraints above.

To model probability theory instead of probabilistic logic, we would change the syntax, but the semantics remain the same. Instead of a propositional language, the agenda would be based on a measurable space as follows: given a measurable space (Ω, F) with Ω the sample space and F the event space, the agenda $X \subseteq \{x_E : E \in F\}$ consists of variables that are indexed by the events in F. Instead of truth values, agents now report probabilities for the variables in X such that a judgment function represents a probability measure for (Ω, F). The Kolmogorov constraints are now constructed with respect to the event space F instead of the language $\mathcal{L}$:

$$\forall x_E \in X: \qquad x_E \geq 0 \qquad \text{(Non-negativity)}$$
$$x_\Omega = 1 \qquad \text{(Normalization)}$$
$$\forall x_{E_1}, x_{E_2} \in X \text{ s.t. } E_1 \cap E_2 = \varnothing: \qquad x_{E_1\cup E_2} = x_{E_1} + x_{E_2} \qquad \text{(Finite additivity)}$$

Yet another alternative for encoding conjunction and disjunction is using constraints with intersection and union. We also change the interpretation for falsum and verum. Let $\Delta_{\mathcal{L},\cap}$ be the following set of constraints:

$$\forall x_\varphi \in X: \qquad x_{\neg\varphi} = \mathcal{T}\backslash x_\varphi \qquad \text{(Negation, difference)}$$
$$\forall x_\varphi, x_\psi \in X: \qquad x_{\varphi\wedge\psi} = x_\varphi \cap x_\psi \qquad \text{(Conjunction, intersection)}$$
$$\forall x_\varphi, x_\psi \in X: \qquad x_{\varphi\vee\psi} = x_\varphi \cup x_\psi \qquad \text{(Disjunction, union)}$$
$$x_\perp = \varnothing \qquad \text{(Falsum set)}$$
$$x_\top = \mathcal{T} \qquad \text{(Verum set)}$$

With this set-theoretic interpretation of the logical operations, we could let the domain constitute the set of possible worlds, or information states, in which agenda-items are true or false. The judgment value J_i then contains the worlds, or states, that agent i considers possible for each variable $x \in X$. By constructing additional constraints, it would be possible to model aggregation problems involving intuitionistic logic, possibly even modal logic.

Example 4.6. In the table below, each row represents an example judgment function that is logically consistent with $\Delta_{\mathcal{L},\cap}$ for the set of truth values $\mathcal{T}$.

$\mathcal{T}$	x_p	x_q	$x_{\neg p}$	$x_{p\wedge q}$	$x_{p\vee q}$
$\{0,1\}$	$\{0,1\}$	$\{1\}$	$\varnothing$	$\{1\}$	$\{0,1\}$
$[0,1]$	$\{0,\frac{3}{4}\}$	$\{\frac{1}{4},1\}$	$(0,\frac{3}{4})\cup(\frac{3}{4},1]$	$\varnothing$	$\{0,\frac{1}{4},\frac{3}{4},1\}$
$[0,1]$	$[0,\frac{2}{3}]$	$[\frac{1}{3},1]$	$(\frac{2}{3},1]$	$[\frac{1}{3},\frac{2}{3}]$	$[0,1]$

4.5 Safety for Profiles

It seems that the proof of Proposition 4.2 is not based on the meaning of the constraint, but rather on the implications it has for the profile structure. We call this structure *unidimensional alignment* (List, 2003). A profile is *unidimensionally aligned* if we can order the agents in such a way that, for each variable, its judgment values can be ordered along the same dimension; for example, each variable's judgment values are in ascending or descending order. Any constraint that imposes a unidimensional alignment of the profile is safe for the median rule, for an odd number of agents, that is.

Proposition 4.5. *For an odd number of agents, given a consistent unidimensionally aligned profile, the median rule will return a consistent outcome.*

Proof. Let $\boldsymbol{J} \in \mathcal{J}(X)^n$ be a consistent unidimensionally aligned profile with odd n. There is a middle agent such that for all $x \in X$, the median judgment value is $J_i(x)$. By definition of the median rule, the collective judgment is J_i and thus consistent. □

Unfortunately, we have this result only for an odd number of agents. For an even number of agents, the (upper) median judgment values might not belong to the same agent, because there is no unique middle agent. Depending on the constraints in place, this could mean that the collective judgment is no longer consistent. To ensure safety for an even number of agents, we would need to change the definition of the median rule or consider other profile structures. The proof of Proposition 4.3 demonstrates how imposing additional constraints can bring about a different profile structure. The structure produced there is called *unidirectional alignment*. A profile is *unidirectionally aligned* if we can order the agents in such a way that, for all variables, their judgment values are in the same ordinal direction; for example, all variables' judgment values are in ascending order. In this way, the median judgment values will always belong to the same agent, thus guaranteeing the consistency of the collective judgment.

Proposition 4.6. *Given a consistent unidirectionally aligned profile, the median rule will return a consistent outcome.*

Proof. Let $\boldsymbol{J} \in \mathcal{J}(X)^n$ be a consistent unidirectionally aligned profile. There is an agent $i \in N$ such that for all $x \in X$, the median judgment value is $J_i(x)$. By definition of the median rule, the collective judgment is J_i and thus consistent. □

List (2003) shows that, in the context of JA, the majority rule always returns a consistent outcome for any unidimensionally aligned profile. We have now translated these ideas into NJA for the median rule. However, the question remains whether we can construct integrity constraints that encode profile structures like unidimensional and unidirectional alignment. Furthermore, for constraints in general, it should be investigated whether specific classes of constraints are safe for certain families of aggregation rules. For example, what class of constraints is safe for bound-quota rules? Determining which constraints are safe for a certain rule can also help define the property that characterizes that rule or family of rules.

Chapter 5

Model Applications

In this chapter, we illustrate the versatility of the NJA framework by presenting some applications to other areas within or related to social choice. These applications highlight the expressiveness and flexibility of NJA and demonstrate how it can unify various aggregation approaches under a common numerical framework. They also serve to motivate further theoretical developments and practical uses of NJA beyond social choice theory, such as public policymaking and multi-agent systems.

In the introduction (Chapter 1), we mentioned the many-valued logic framework of judgment aggregation by Pauly and van Hees (2006) and the probabilistic framework of opinion pooling. NJA can model both of these approaches. By treating valuations of many-valued logic as point-valued judgments within the range of the possible truth values, we can capture graded acceptance and rejection. In addition to restrictions on the input, we can incorporate constraints, such as those presented in Section 4.4, to encode the logical dependency between agenda-items of a many-valued JA problem. Similarly, we can model the probabilistic framework of opinion pooling as described by Elkin and Pettigrew (2025). By treating credences and imprecise probabilities as, respectively, point-valued and set-valued judgments ranging between 0 and 1, NJA accommodates opinion pooling methods, such as linear pooling and multiplicative pooling, as well as imprecise probability pooling. In addition, integrity constraints can be implemented to encode the probabilistic dependency between agenda-items of an opinion pooling problem. We can thus apply NJA to develop constraint-based versions of traditionally formula-based aggregation frameworks.

We will now discuss some other applications in more detail. First, we discuss how preference aggregation can be embedded within NJA via standard judgment aggregation and hint towards a direct embedding by interpreting preferences as numerical rankings. Second, we demonstrate how binary JA can be modeled as a special case of NJA using integrity constraints to encode certain judgment properties. Lastly, we present the preliminaries of a model for the aggregation of societal trade-offs with integrity constraints to maintain their structural consistency.

5.1 Preference Aggregation

List and Pettit (2004) were the first to model preference aggregation problems as special cases of judgment aggregation problems. They use predicate logic to represent a preference ordering as a set of binary preference judgments; a first-order formula xPy, where x and y are issues of the preference aggregation problem, expresses "x is strictly preferred to y". These first-order formulas are then used to form the agenda of the corresponding JA problem. Dietrich and List (2007a) build on this idea and construct an explicit embedding of preference aggregation into binary judgment aggregation. Endriss (2018) adapts this approach by incorporating integrity constraints to encode the structural properties of preference relations. In the next section, we will see how binary JA can be modeled within NJA. Combined with this work, we must also be able to model preference aggregation directly within NJA.

Grandi and Endriss (2011) extend the connection between preference aggregation and judgment aggregation by showing how both can be treated within a unified binary aggregation framework with integrity constraints to enforce structural properties.

Grossi (2009) takes a different approach outside the binary setting. By borrowing ideas from logical semantics, he shows how preference aggregation can be considered a special case of judgment aggregation in a many-valued logic setting, but also the other way around, claiming that they can be studied as two faces of the same coin. Back in 1954, Debreu already proved that every complete, transitive, and continuous preference relation can be represented by a continuous ordinal utility function.[1] Grossi (2009) uses this result to define ranking functions that attribute a ranking value to agenda-items of the preference aggregation problem. Representing preferences as numerical ranking values can be directly adopted in the NJA framework. This makes the direct embedding from preference aggregation into NJA even more intuitive than the one Grossi presents. For the same reason, an NJA embedding using Debreu's Theorem is more natural than one via binary (judgment) aggregation because it incorporates the numerical nature of the framework as a whole.

5.2 Binary Judgment Aggregation

Binary JA incorporates an agenda that contains compound logical formulas, such as $p \wedge q$. In contrast, binary aggregation and NJA do not allow complex variables. However, Grandi and Endriss (2011) show how binary aggregation can embed binary JA using integrity constraints. For a complex agenda Φ, the general agenda contains binary variables indexed by the formulas in Φ and their negations. Then, integrity constraints are used to encode the logical semantics of the indices of the variables, i.e., the original formulas in Φ. In Section 4.4, we have seen how integrity constraints can be constructed to create logical dependencies between variables. So, if we can translate the binary aggregation model and the associated integrity constraints into NJA, we infer that binary JA can also be embedded in NJA. Let us see what this would look like.

Firstly, we translate the propositional agenda Φ into a general agenda X^Φ such that $X^\Phi = \{x_\varphi : \varphi \in \Phi\}$. Secondly, we only allow for bivalent judgment functions in the input profile $\boldsymbol{J} \in \mathcal{J}(X^\Phi)^n$. This means that every variable $x_\varphi \in X^\Phi$ receives the (truth) judgment value $J(x_\varphi) = \mathbf{1}$ if φ is supported and $J(x_\varphi) = \mathbf{0}$ if it is rejected. We observe that a bivalent judgment function $J \in \mathcal{J}(X^\Phi)$ now emulates a judgment set S (in the interpretation of JA) on the propositional agenda Φ: for each formula $\varphi \in \Phi$, $\varphi \in S$ is equivalent to $J(x_\varphi) = \mathbf{1}$ and $\varphi \notin S$ is equivalent to $J(x_\varphi) = \mathbf{0}$. Because bivalent judgments are by definition ranged (see Definition 2.1), we can apply any of the judgment aggregation rules introduced in Section 2.2. However, the output of these rules might differ given bivalent input as opposed to ranged input.

Secondly, two central notions of JA are completeness and complement-freeness (Endriss, 2016). A judgment set S is called *complete* if $\varphi \in S$ or $\sim\varphi \in S$ for all $\varphi \in \Phi$ and S is called *complement-free* if $\varphi \notin S$ or $\sim\varphi \notin S$ for all $\varphi \in \Phi$.[2] We define the following NJA integrity constraints to encode these notions.

[1] A continuous function that represents the preferences of an agent on an ordinal scale.

[2] We define $\sim\varphi := \psi$ if $\varphi = \neg\psi$ and $\sim\varphi := \neg\varphi$ otherwise.

For all positive $\varphi \in \Phi$,

$$x_\varphi + x_{\sim\varphi} \geq 1 \qquad \text{(completeness)}$$
$$x_\varphi + x_{\sim\varphi} \leq 1 \qquad \text{(complement-freeness)}$$

Consider a JA profile and translate it into a bivalent profile $\boldsymbol{J} \in \mathcal{J}(X^\Phi)^n$ such that for every agent i, we let $J_i(x_\varphi) = \mathbf{1}$ if they support φ in the JA profile and $J(\varphi) = \mathbf{0}$ if they reject it. It is easy to check that $\boldsymbol{J}$ satisfies the completeness constraint if and only if the original JA profile is complete, and $\boldsymbol{J}$ satisfies the complement-freeness constraint if and only if the original JA profile is complement-free.

Let us examine the safety of these constraints for the majority support rule (2.10). On bivalent profiles, completeness and complement-freeness induce a unidimensional alignment of the profile. This means that the middle judgment represents the judgment of the majority of agents. This judgment is clearly complete or complement-free if the profile is complete or complement-free, respectively. However, there is only a middle judgment with an odd number of agents. If the number is even and there is no strict majority in favor of either (truth) value, then the majority support rule will pick neither, resulting in a complement-free yet incomplete collective judgment.

To conclude, we point out that the majority support rule (2.10) is the NJA generalization of the (strict) majority rule from JA. The equivalence is quite obvious. Firstly, their definitions work equivalently on bivalent input. The JA majority rule is defined as follows. It takes a profile of judgment sets and returns the set of propositions that are supported by more than half of the agents. Let us refer to the majority rules as F^{NJA} and F^{JA}, respectively. Clearly, for any proposition $\varphi \in \Phi$, we have $F^{\text{NJA}}(\boldsymbol{J})(x_\varphi) = \mathbf{1}$ if $\varphi \in F^{\text{JA}}(\boldsymbol{J})$ and $F^{\text{NJA}}(\boldsymbol{J})(x_\varphi) = \mathbf{0}$ if $\varphi \notin F^{\text{JA}}(\boldsymbol{J})$.

Furthermore, their properties are similarly correlated. For an odd number of agents, the JA majority rule always returns a complete and complement-free outcome (List & Pettit, 2002). Similarly, the completeness and complement-freeness constraints are safe for the majority support rule, as we have just discussed. Moreover, we know that the majority support rule is anonymous, locally neutral, locally independent, and has monotonic support by the corollaries of Theorem 3.4. These notions have their origin in the corresponding properties of anonymity, neutrality, independence, and monotonicity from JA, which, together with completeness and complement-freeness, uniquely characterize the JA majority rule (Endriss, 2016).

5.3 Societal Trade-Offs

In many public policy contexts, stakeholders must navigate trade-offs between competing objectives, such as reducing pollution, conserving resources, promoting economic growth, and ensuring public health. These goals often conflict, and pursuing one to a greater extent may require compromising on another. The judgment or decision about how much of an activity, such as using gasoline, is considered equivalent in societal cost to another, such as producing trash, is called a societal trade-off. Conitzer et al. (2015) discuss the considerations for a systematic approach to obtain such numerical trade-off values at the societal level based on individual input. They observe that applying the median rule to each trade-off can produce globally inconsistent results. To address this problem, Conitzer et al. (2016) introduce a family of distance-based aggregation rules, where the collective trade-offs minimize the overall disagreement with the individual trade-off values.

In this section, I present a point-valued NJA model to explicitly simulate the aggregation of societal trade-offs as proposed by Conitzer et al. (2016). The fundamental idea is that agents submit numerical trade-off values for pairs of activities, forming a weighted directed graph, and the aim is to derive a globally consistent collective trade-off structure.

Let A be a finite set of activities, such as clearing forest. The agenda, X, now consists of variables indexed by ordered pairs of these activities such that for every pair of activities $a, b \in A$, either $x_{ab} \in X$ or $x_{ba} \in X$, but not both. Furthermore, we require that the input profiles be positive and point-valued rather than ranged. This means that every variable $x \in X$ receives a judged trade-off value $J(x) = \boldsymbol{k}$ for $k \in \mathbb{R}_{>0}$. A trade-off judgment function J is consistent if it satisfies the following integrity constraints:

$$\forall x_{ab}, x_{bc}, x_{ac} \in X : x_{ab} \times x_{bc} = x_{ac} \tag{5.1}$$

$$\forall x_{ab}, x_{bc}, x_{ca} \in X : x_{ab} \times x_{bc} = 1/x_{ca} \tag{5.2}$$

A trade-off judgment $J \in \mathcal{J}(X)$ can be represented as the weight function of a weighted directed graph, as illustrated in the figure below.

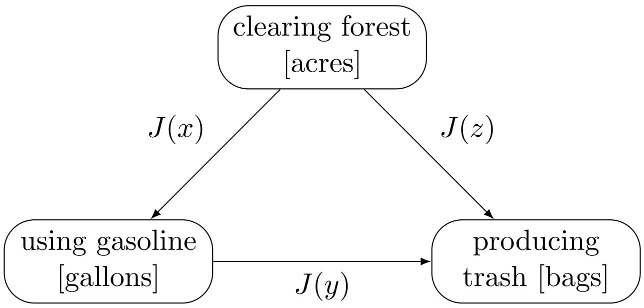

Figure 5.1: A weighted directed graph representation of a societal trade-off judgment. A directed edge $x_{ab} \in X$ between two nodes $a, b \in A$ with weight $J(x_{ab})$ represents that one unit of activity a is considered equivalent to $J(x_{ab})$-many units of activity b. The trade-off in the graph above is consistent if $J(x) \times J(y) = J(z)$.

Conitzer et al. (2015) bring up the issue that applying the median rule to individually consistent trade-offs can lead to inconsistencies in the collective structure. Indeed, as we have seen in Example 4.2 of the previous chapter, multiplication with three variables is not safe for the median rule. The figure below illustrates a simple example of this paradox for societal trade-offs.

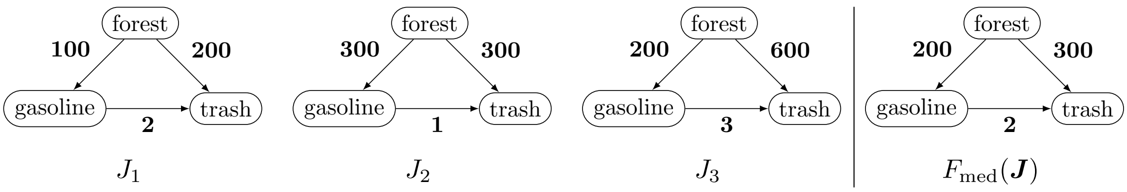

Figure 5.2: An example of the doctrinal paradox for societal trade-offs. Although the left three graphs are consistent, applying the median rule on the profile results in the rightmost graph; the inconsistent collective trade-off structure with $\mathbf{300} \neq \mathbf{2} \times \mathbf{200}$.

In search of another aggregation rule that guarantees consistent trade-offs, Conitzer et al. (2016) introduce a family of distance-based aggregation rules that aggregate judgments by minimizing total disagreement. Disagreement is measured using a distance function, e.g., linear distance or logarithmic distance. What makes distance-based rules appealing is that they reduce to the median rule when aggregating a single trade-off (that is, involving only two activities), and they can be interpreted as maximum likelihood estimators of the truth.

To axiomatically characterize this family of rules, they introduce two properties: independence of choice of units (ICU) and independence of edge directions (IED). An aggregation rule satisfies ICU if the outcome does not depend on the units in which activities are measured. For example, it should not matter whether we measure the forest area in acres or hectares, or whether we measure the gasoline usage in gallons or liters. An aggregation rule satisfies IED if the outcome does not depend on the edge direction between two activities. That is, it should not matter whether the trade-off between two activities a and b is represented by the variable x_{ab} or x_{ba}.

Conitzer et al. (2016) illustrate how the linear distance-based rule, a variant of the Kemeny rule from preference aggregation, violates ICU and IED because larger values carry more weight in the distance function. On the other hand, the logarithmic distance-based rule satisfies both properties. Moreover, they show that outcomes under the logarithmic distance-based rule can be computed in polynomial time. This makes for a good candidate to aggregate societal trade-offs.

To conclude, Conitzer et al. (2016) raise the concern that agents might be more comfortable reporting an interval for each trade-off between activities. We can handle this in the NJA framework by aggregating the supremum and infimum separately. In addition, it is straightforward to translate the definitions of distance-based aggregation rules and the properties of ICU and IED into our NJA framework for point-valued judgments. However, the question remains how to extend these rules and properties for intervals and how to create a trade-off structure based on intervals.

Chapter 6

Discussion and Conclusion

This thesis has proposed and explored a general framework for numerical judgment aggregation (NJA), motivated by the limitations of classical aggregation models. NJA provides a flexible and expressive approach to aggregating numerical judgments, accommodating a broader spectrum of epistemic and practical scenarios. In this concluding chapter, we summarize the main contributions, reflect on key insights and limitations, and outline directions for future research.

We began by defining the core framework in Chapter 2, introducing agenda-items as variables that take judgments in the form of sets of reals and allow agents to express uncertainty or indifference through intervals or sets. We defined several classes of aggregation rules, including calculation-based rules (e.g., the mean rule), ordinal rules (e.g., the median rule), and support-based rules (e.g., the majority support rule), and we defined repair operations to transform outputs into more manageable or appropriate forms, such as intervals or points.

In Chapter 3, we investigated the axiomatic properties of these rules. The axioms, such as anonymity, neutrality, autonomy, and monotonic support, allowed us to evaluate the normative behavior of the aggregation rules. A key insight was that the family of support-quota rules can be axiomatically characterized using the concept of winning coalitions, providing a natural correspondence with characterization results of quota rules in standard judgment aggregation.

Chapter 4 defined the formal notions for the use of integrity constraints for NJA and provided some results on the consistency of judgments with respect to certain constraints. We distinguished between different classes of constraints (e.g., rationality and feasibility, and arithmetic and algebraic) and showed how constraints can be used to encode consistency, coherence, and interdependence across different agendas.

Finally, in Chapter 5, we illustrated the practical scope of NJA through applications for preference aggregation, binary judgment aggregation, and the modeling of societal trade-offs. These examples demonstrated that NJA can simulate classical aggregation frameworks as well as offer ways to analyze concepts in areas requiring richer numerical judgments.

NJA contributes to the literature by bridging the gaps between preference aggregation, formula-based judgment aggregation, and probabilistic opinion pooling. Its flexibility allows it to unify various aggregation domains while preserving structure through integrity constraints. The most unique feature of NJA, as opposed to other aggregation frameworks, is that it accommodates not only bivalent or point-valued judgments, but also ranged and set-valued input. We have focused mainly on closed intervals (including point-valued judgments); incorporating open intervals or arbitrary sets requires further formal development. Fortunately, the extension seems natural for most definitions and will probably open the door for more characterization and impossibility results.

At the same time, the generality of NJA raises new challenges. The expressiveness, for example, is heavily dependent on the constraints that need to be constructed and implemented for specific applications. For example, many rules require a certain input structure. In contrast, other rules require repair operations to yield manageable output. Unfortunately, the benefits of applying a repair operation can be at the expense of the normative appeal of the aggregation procedure. Lastly, we have no formal tool yet to axiomatically characterize aggregation rules that cannot be defined by winning coalitions, i.e., rules that are not locally independent.

Looking ahead, several directions for further research emerge. First, any practical implementation of NJA rules, especially under integrity constraints, raises questions about computational complexity and feasibility, even more so than for binary (judgment) aggregation because of the real-valued judgments. Second, extending the input format even further, e.g., incorporating judgments in the form of open intervals, arbitrary sets, or even empty sets, could expand the framework's applicability. However, many rules, axiomatic properties, and integrity constraints have not been well-defined on such judgments. For example, the infimum of an open set is not contained in the set itself, and the supremum of an unbounded set can be infinity. What meaning do they carry then, and what are the implications for the rest of the aggregation procedure?

Another research direction concerns the behavior of agents. A pragmatic question is how to deal with outliers and extreme values, but it is also connected to the need for a theoretical analysis of strategic manipulation for NJA. Although some observations have been made throughout this work, no systematic study has yet been conducted. We would need to define measures that describe how satisfied agents are with the outcome and whether they have an incentive to lie about their true judgment.

Lastly, we hope that the framework can be implemented to model applied aggregation systems to provide theoretical insights into real-world challenges. Rey et al. (2020) have already shown that their formal framework for participatory budgeting, a variant of fair allocation with a budget constraint on the outcome, can be embedded in binary aggregation with integrity constraints. Given that NJA incorporates numerical judgments and integrity constraints, we can embed participatory budgeting directly into NJA. There are probably many more aggregation systems related to the area of social choice that can be captured by NJA.

In conclusion, this thesis has laid the foundations for a framework of numerical judgment aggregation; a principled, versatile, and extensible framework for collective decision-making that requires more nuance or deals with uncertainty. We hope that it provides new insights and inspiration for further exploration at the intersection of logic and social choice.

Appendix A

Sets, Intervals, and Operations

A.1 Sets and Intervals

Notation Depending on readability, the cardinality of a set S is denoted by either $\#S$ or $|S|$. Sets of numbers are denoted by a blackboard bold letter, such as $\mathbb{R}$. Collections or families of sets are usually denoted by a calligraphic letter, such as $\mathcal{J}$. Tuples are written in boldface, such as $\boldsymbol{J} = (J_1, \ldots, J_n)$. If a tuple has the subscript $-i$, it denotes the tuple without the element associated with i. For example, $\boldsymbol{J}_{-i} = (J_1, \ldots, J_{i-1}, J_{i+1}, \ldots, J_n)$. The set of functions from A to B is denoted B^A.

Upper and lower bounds A set $S \subset \mathbb{R}$ of real numbers is bounded from above if there exists a number $b \in \mathbb{R}$, such that $b \geq k$ for every $k \in S$; b is called an *upper bound* of S. Similarly, S is bounded from below if there exists a number $a \in \mathbb{R}$, such that $a \leq k$ for every $k \in S$; a is called a *lower bound* of S. A set is called *bounded* if it is bounded from both above and below.

Supremum and infimum If $b \in \mathbb{R}$ is an upper bound of S such that $b \leq b'$ for every upper bound b' of S in $\mathbb{R}$, then b is called the *supremum* of S, denoted $\sup S$. The supremum of S is its least upper bound. Similarly, if $a \in \mathbb{R}$ is a lower bound of S such that $a \geq a'$ for every lower bound a' of S in $\mathbb{R}$, then a is called the *infimum* of S, denoted $\sup X$. The infimum of S is its greatest lower bound.

If S is not bounded from above, then the supremum does not exist, and we write $\sup X = \infty$. If S is not bounded from below, then the infimum does not exist, and we write $\inf S = -\infty$. The supremum and infimum of a set S are unique if they exist, and if both exist, then $\inf S \leq \sup A$. If $S = \varnothing$, is the empty set, then every real number is both an upper and a lower bound of S, and we write $\sup S = -\infty$ and $\inf S = \infty$. If S is non-empty, then $\sup S \geq \inf S$. If S and T are both sets of real numbers such that $S \subseteq T$, then $\inf S \geq \inf T$ and $\sup S \leq \sup T$.

Open and closed sets The infimum of a set S does not necessarily belong to S itself. If it does, that is, $\inf S \in S$, it is also called a minimum of S. Similarly, if the supremum of S is an element of S, i.e., $\sup S \in S$, it is called a maximum of S. A set is closed if it has both a minimum and a maximum, it is open if it has neither, and it is half-open if it has only one.

Convex sets A set of points is *convex* if it contains every line segment between two points in the set.

Real intervals An interval is a convex subset of $\mathbb{R}$. The set of all reals $\mathbb{R}$ can be given by the open interval $(-\infty, \infty)$. Intervals can be closed, e.g.,$[0, 100]$; open, e.g., $(-1, 1)$; or half-open, e.g., $[19, \infty)$; and bounded, e.g., $(-459, 32)$, or not, e.g., $(-\infty, \infty)$. A singleton $\{a\}$ of the real number a can be written as the interval $[a, a]$. The empty set can be given by any interval of the form (a, a), $(a, a]$, or $[a, a)$.

A.2 Set Operations

Intersection The *intersection* of two sets S and T, denoted $S \cap T$, is the set of elements that are in both S and T:

$$S \cap T = \{x : x \in S \text{ and } x \in T\}$$

The intersection of two intervals, although possibly empty, is always an interval.

Union The *union* of two sets S and T, denoted $S \cup T$, is the set of elements that are in S or T, or both:

$$S \cup T = \{x : x \in S \text{ or } x \in T\}$$

If S and T are *disjoint*, i.e., $S \cap T = \varnothing$, the *disjoint union*, denoted $S \uplus T$, is the set of elements that are in S or T.

The union of two intervals is an interval if and only if it is convex, meaning there are no 'gaps'. That is, either they have a non-empty intersection or an open end-point of one interval is a closed end-point of the other, for example $(a, b) \cup [b, c] = (a, c]$.

Difference The *(set) difference* between two sets S and T, denoted $S \setminus T$, is the set of elements that are in S but not in T:

$$S \setminus T = \{x \in S : x \notin T\}$$

$S \setminus T$ is called the *complement* of T in S; it is denoted $\overline{T}$ whenever S is clear from context. The complement of a closed set in $\mathbb{R}$ is open.

Convex hull The *convex hull* of a set S, denoted $\text{hull}\,(S)$, is the smallest convex set that contains S. The convex hull of a real interval is the interval itself. The convex hull of an arbitrary set of real numbers $S \subseteq \mathbb{R}$ is the smallest interval I that contains S:

$$\text{hull}\,(S) = \arg\min_{I \supseteq S} |I|$$

Convex core The *convex core* of a set S, denoted $\text{core}\,(S)$, is the largest convex set that S contains. Note that the convex core might not be unique. The convex core of a set of real numbers $S \subseteq \mathbb{R}$ is the largest interval I (possibly a point) that it contains:

$$\text{core}\,(S) \in \arg\max_{I \subseteq S} |I|$$

Multiplication The *product* of two sets S and T, denoted $S \times T$, is the set of all pairwise products:

$$S \times T = \{x \times y : x \in S,\ y \in T\}$$

The *linear transformation* of a set S by a factor $k \in \mathbb{R}$, denoted $k \times S$, is a special case of setwise multiplication, where k is treated as a singleton $\{k\}$:

$$k \times S = \{k \times x : x \in S\}$$

For closed intervals, multiplication simplifies to:

$$\begin{aligned}[a, b] \times [a', b'] = \Big[&\min \{a \times a', a \times b', b \times a', b \times b'\}, \\ &\max \{a \times a', a \times b', b \times a', b \times b'\} \Big]\end{aligned}$$

Addition The *(Minkowski) sum* of two sets S and T, denoted $S + T$, is the set of all pairwise sums:

$$S + T = \{x + y : x \in S,\ y \in T\}$$

The *translation* of a set S by a term $k \in \mathbb{R}$, denoted $S + k$, is a special case of setwise addition, where k is treated as a singleton $\{k\}$:

$$S + k = \{x + k : x \in S\}$$

For closed intervals, addition simplifies to:

$$[a, b] + [a', b'] = [a + a', b + b']$$

Subtraction Using setwise multiplication and addition, we can calculate the *(arithmetic) difference* between two sets S and T, denoted $S - T$, as the set of all pairwise differences:

$$\begin{aligned} S - T &= S + (-T) \\ &= S + (-1 \times T) \\ &= \{x - y : x \in S,\ y \in T\} \end{aligned}$$

Probabilistic addition The *probabilistic sum* of two sets S and T, denoted $S \mathbin{\hat{+}} T$, is the set of all pairwise sums normalized with the pairwise product:

$$S \mathbin{\hat{+}} T = \{x + y - x \times y : x \in S,\ y \in T\}$$

Maximum and minimum The *maximal elements* of two sets S and T, denoted $\max(S, T)$, is the set of all pairwise maxima:

$$\max(S, T) = \{\max(x, y) : x \in S,\ y \in T\}$$

Analogously, the *minimal elements* of two sets S and T, denoted $\min(S, T)$, is the set of all pairwise minima:

$$\min(S, T) = \{\min(x, y) : x \in S,\ y \in T\}$$

Appendix B

Proof of Lemma 3.3

Lemma 3.3 (Winning Coalitions)**.** *Let F be a locally independent aggregation rule, and, for every $x \in X$ and $k \in \mathbb{R}$, let $\mathcal{W}_{x,k} \subseteq 2^N$ be the corresponding family of winning coalitions. Then we have the following:*

(i) F is anonymous if and only if $\mathcal{W}_{x,k}$ is closed under equinumerosity; $C \in \mathcal{W}_{x,k}$ and $|C| = |C'|$ imply $C' \in \mathcal{W}_{x,k}$ for all $C, C' \subseteq N$ and all $x \in X$ and $k \in \mathbb{R}$.

(ii) F is a dictatorship if and only if there exists an $i \in N$ such that for all $x \in X$ and $k \in \mathbb{R}$, $C \in \mathcal{W}_{x,k} \Leftrightarrow i \in C$.

(iii) F is unanimously supporting if and only if $N \in \mathcal{W}_{x,k}$ for all $x \in X$ and $k \in \mathbb{R}$.

(iv) F is unanimously refuting if and only if $\varnothing \notin \mathcal{W}_{x,k}$ for all $x \in X$ and $k \in \mathbb{R}$.

(v) F is locally neutral if and only if $\mathcal{W}_{x,k} = \mathcal{W}_{y,k}$ for all $x, y \in X$ and $k \in \mathbb{R}$.

(vi) F is autonomous if and only if $\mathcal{W}_{x,k} = \mathcal{W}_{x,k'}$ for all $x \in X$ and $k, k' \in \mathbb{R}$.

(vii) F has monotonic support if and only if $\mathcal{W}_{x,k}$ is upward-closed; $C \in \mathcal{W}_{x,k}$ and $C \subseteq C'$ imply $C' \in \mathcal{W}_{x,k}$ for all $C, C' \subseteq N$, $x \in X$ and $k \in \mathbb{R}$.

Proof. We prove the claims for each property separately.

(i) ($\Rightarrow$) Suppose that F is anonymous. Consider a profile $\boldsymbol{J} \in \mathcal{J}(X)^n$ and a coalition $C \subseteq N$ such that $N^{\boldsymbol{J}}_{x,k} \in \mathcal{W}_{x,k}$ and $|N^{\boldsymbol{J}}_{x,k}| = |C|$. To show that $\mathcal{W}_{x,k}$ is closed under equinumerosity, we have to show that $C \in \mathcal{W}_{x,k}$.

Let $\pi : N \to N$ be a permutation that maps every agent in $N^{\boldsymbol{J}}_{x,k}$ to an agent in C, i.e., $\pi(N^{\boldsymbol{J}}_{x,k}) = C$. It follows from the anonymity of F that $F(J_1, \ldots, J_n) = F(J_{\pi(1)}, \ldots, J_{\pi(n)})$. This means that C must also be winning, i.e., $C \in \mathcal{W}_{x,k}$.

($\Leftarrow$) Suppose $\mathcal{W}_{x,k}$ is closed under equinumerosity. Consider a profile $\boldsymbol{J} \in \mathcal{J}(X)^n$ and a permutation $\pi : N \to N$. Let $\boldsymbol{J}^\pi$ be the shorthand for the profile $(J_{\pi(1)}, \ldots, J_{\pi(n)})$. To show that F is anonymous, we have to show that $F(\boldsymbol{J}) = F(\boldsymbol{J}^\pi)$.

Obviously, $N^{\boldsymbol{J}}_{x,k} = N^{\boldsymbol{J}^\pi}_{x,k}$, so we have $N^{\boldsymbol{J}}_{x,k} \in \mathcal{W}_{x,k}$ if and only if $N^{\boldsymbol{J}^\pi}_{x,k} \in \mathcal{W}_{x,k}$. This means that $k \in F(\boldsymbol{J})(x)$ if and only if $k \in F(\boldsymbol{J}^\pi)(x)$. Because this holds for all x and k, we have $F(\boldsymbol{J}) = F(\boldsymbol{J}^\pi)$.

(ii) This part of the proof works both ways simultaneously:

$$
\begin{aligned}
&F \text{ is a dictatorship}\\
&\Leftrightarrow \qquad \text{(def. of dictatorship)}\\
&F(\boldsymbol{J}) = J_i\\
&\Leftrightarrow \qquad \text{(for all } x \text{ and } k)\\
&k \in F(\boldsymbol{J})(x) \text{ iff } k \in J_i(x)\\
&\Leftrightarrow \qquad \text{(def. of support)}\\
&k \in F(\boldsymbol{J})(x) \text{ iff } i \in N^{\boldsymbol{J}}_{x,k}\\
&\Leftrightarrow \qquad \text{(def. of winning coalition)}\\
&C \in \mathcal{W}_{x,k} \text{ iff } i \in C
\end{aligned}
$$

(iii) ($\Rightarrow$) Suppose F is unanimously supporting. This means that $J_1(x) \cap \cdots \cap J_n(x) \subseteq F(\boldsymbol{J})(x)$ for all $x \in X$. Clearly, then, N must be a winning coalition for all $x \in X$ and $k \in \mathbb{R}$.

($\Leftarrow$) Suppose N is a winning coalition for all $x \in X$ and $k \in \mathbb{R}$. If $k \in J_1(x) \cap \cdots \cap J_n(x)$, we have $N^{\boldsymbol{J}}_{x,k} = N$. Because N is winning, $k \in F(\boldsymbol{J})(x)$. This reasoning works for all k, so $J_1(x) \cap \cdots \cap J_n(x) \subseteq F(\boldsymbol{J})(x)$, i.e., F is unanimously supporting.

(iv) ($\Rightarrow$) Suppose F is unanimously refuting. This means that $J_1(x) \cup \cdots \cup J_n(x) \supseteq F(\boldsymbol{J})(x)$ for all $x \in X$. Clearly, then, the empty set cannot be a winning coalition for any $x \in X$ and $k \in \mathbb{R}$.

($\Leftarrow$) Suppose the empty set is not a winning coalition for any $x \in X$ and $k \in \mathbb{R}$. If $k \in F(\boldsymbol{J})(x)$, then it must be that $N^{\boldsymbol{J}}_{x,k} \neq \varnothing$. This means there must be some $i \in N$ such that $k \in J_i(x)$, i.e., $k \in J_1(x) \cup \cdots \cup J_n(x)$. This reasoning works for all k, so $J_1(x) \cup \cdots \cup J_n(x) \supseteq F(\boldsymbol{J})(x)$, i.e., F is unanimously refuting.

(v) ($\Rightarrow$) Suppose that F is locally neutral. We show that $\mathcal{W}_{x,k} = \mathcal{W}_{y,k}$. Neutrality states that $k \in J_i(x) \Leftrightarrow k \in J_i(y)$ for all $i \in N$ implies $k \in F(\boldsymbol{J})(x) \Leftrightarrow k \in F(\boldsymbol{J})(y)$. This means that $N^{\boldsymbol{J}}_{x,k} = N^{\boldsymbol{J}}_{y,k}$ implies $N^{\boldsymbol{J}}_{x,k} \in \mathcal{W}_{x,k} \Leftrightarrow N^{\boldsymbol{J}}_{y,k} \in \mathcal{W}_{y,k}$. Since this holds for all profiles, it must be that $\mathcal{W}_{x,k} = \mathcal{W}_{y,k}$.

($\Leftarrow$) Assume that $\mathcal{W}_{x,k} = \mathcal{W}_{y,k}$. If we have $k \in J_i(x) \Leftrightarrow k \in J_i(y)$ for all $i \in N$, then $N^{\boldsymbol{J}}_{x,k} = N^{\boldsymbol{J}}_{y,k}$. By our assumption, it follows that $N^{\boldsymbol{J}}_{x,k} \in \mathcal{W}_{x,k} \Leftrightarrow N^{\boldsymbol{J}}_{y,k} \in \mathcal{W}_{y,k}$. Thus, $k \in J_i(x) \Leftrightarrow k \in J_i(y)$ for all $i \in N$ implies $k \in F(\boldsymbol{J})(x) \Leftrightarrow k \in F(\boldsymbol{J})(y)$, i.e., F is locally neutral.

(vi) ($\Rightarrow$) Suppose that F is autonomous. We show that $\mathcal{W}_{x,k} = \mathcal{W}_{x,k'}$. Autonomy states that $k \in J_i(x) \Leftrightarrow k' \in J_i(x)$ for all $i \in N$ implies $k \in F(\boldsymbol{J})(x) \Leftrightarrow k' \in F(\boldsymbol{J})(x)$. This means that $N^{\boldsymbol{J}}_{x,k} = N^{\boldsymbol{J}}_{x,k'}$ implies $N^{\boldsymbol{J}}_{x,k} \in \mathcal{W}_{x,k} \Leftrightarrow N^{\boldsymbol{J}}_{x,k'} \in \mathcal{W}_{x,k'}$. Since this holds for all profiles, it must be that $\mathcal{W}_{x,k} = \mathcal{W}_{x,k'}$.

($\Leftarrow$) Assume that $\mathcal{W}_{x,k} = \mathcal{W}_{x,k'}$. If we have $k \in J_i(x) \Leftrightarrow k' \in J_i(x)$ for all $i \in N$, then $N^{\boldsymbol{J}}_{x,k} = N^{\boldsymbol{J}}_{x,k'}$. By our assumption, it follows that $N^{\boldsymbol{J}}_{x,k} \in \mathcal{W}_{x,k} \Leftrightarrow N^{\boldsymbol{J}}_{x,k'} \in \mathcal{W}_{x,k'}$. Thus, $k \in J_i(x) \Leftrightarrow k' \in J_i(x)$ for all $i \in N$ implies $k \in F(\boldsymbol{J})(x) \Leftrightarrow k' \in F(\boldsymbol{J})(x)$, i.e., F is autonomous.

(vii) ($\Rightarrow$) Suppose that F has monotonic support. Consider a profile $\boldsymbol{J} \in \mathcal{J}(X)^n$ and a coalition $C \subseteq N$ such that $N^{\boldsymbol{J}}_{x,k} \in \mathcal{W}_{x,k}$ and $N^{\boldsymbol{J}}_{x,k} \subseteq C$. To show that $\mathcal{W}_{x,k}$ is upward-closed, we have to show that $C \in \mathcal{W}_{x,k}$.

For all $i \in C \backslash N^{\boldsymbol{J}}_{x,k}$, let $J'_i(x) = J_i(x) \cup \{k\}$. Consider some $i \in C \backslash N^{\boldsymbol{J}}_{x,k}$. By definition of monotonic support, it follows that $k \in F(\boldsymbol{J}_{-i}, J'_i)(x)$. This means that $N^{\boldsymbol{J}}_{x,k} \cup \{i\} \in \mathcal{W}_{x,k}$. If we repeat this consecutively for all $i \in C \backslash N^{\boldsymbol{J}}_{x,k}$, we get $N^{\boldsymbol{J}}_{x,k} \cup (C \backslash N^{\boldsymbol{J}}_{x,k}) \in \mathcal{W}_{x,k}$, i.e., $C \in \mathcal{W}_{x,k}$.

($\Leftarrow$) Assume that $\mathcal{W}_{x,k}$ is upward-closed. Consider a profile $\boldsymbol{J} \in \mathcal{J}(X)^n$ and a judgment $J'_i \in \mathcal{J}(X)$ such that $k \in J'_i(x) \backslash J_i(x)$ and $J'_i(y) = J_i(y)$ for all $y \in X$ such that $y \neq x$, and $k \in F(\boldsymbol{J})(x)$. Because $k \in F(\boldsymbol{J})(x)$, we have $N^{\boldsymbol{J}}_{x,k} \in \mathcal{W}_{x,k}$, and because $k \in J'_i(x) \backslash J_i(x)$, we have $N^{\boldsymbol{J}}_{x,k} \subseteq N^{(\boldsymbol{J}_{-i}, J'_i)}_{x,k}$. By our assumption, it follows that $N^{(\boldsymbol{J}_{-i}, J'_i)}_{x,k} \in \mathcal{W}_{x,k}$, which means that $k \in F(\boldsymbol{J}_{-i}, J'_i)(x)$. Thus, F has monotonic support. $\square$

Bibliography

Arrow, K. J. (1950). A Difficulty in the Concept of Social Welfare. *The Journal of Political Economy*, *58*(4), 328–346.

Condorcet, N. d. (1785). *Essai sur l'Application de l'Analyse à la Probabilité des Décisions Rendues à la Pluralité des Voix*. Imprimerie Royale.

Conitzer, V., Brill, M., & Freeman, R. (2015). Crowdsourcing Societal Tradeoffs. *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2015)*, 1213–1217.

Conitzer, V., Freeman, R., Brill, M., & Li, Y. (2016). Rules for Choosing Societal Tradeoffs. *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI-2016)*, 460–467.

Debreu, G. (1954). Representation of a Preference Ordering by a Numerical Function. In M. Thrall, R. C. Davis, & C. H. Coombs (Eds.), *Decision Processes* (pp. 159–165). John Wiley; Sons.

Dietrich, F., & List, C. (2007a). Arrow's Theorem in Judgment Aggregation. *Social Choice and Welfare*, *29*(1), 19–33.

Dietrich, F., & List, C. (2007b). Strategy-Proof Judgment Aggregation. *Economics and Philosophy*, *23*(3), 269–300.

Dietrich, F., & List, C. (2017). Probabilistic Opinion Pooling Generalized. Part One: General Agendas. *Social Choice and Welfare*, *48*, 747–786.

Elkin, L., & Pettigrew, R. (2025). *Opinion Pooling*. Cambridge University Press.

Endriss, U. (2016). Judgment Aggregation. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, & A. D. Procaccia (Eds.), *Handbook of Computational Social Choice* (pp. 399–426). Cambridge University Press.

Endriss, U. (2018). Judgment Aggregation with Rationality and Feasibility Constraints. *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2018)*.

Fan, K. (1959). *Convex Sets and Their Applications*. Argonne National Laboratory.

Gödel, K. (1932). Zum intuitionistischen Aussagenkalkül. *Anzeiger der Akademie der Wissenschaften in Wien*, *69*, 65–66.

Grandi, U., & Endriss, U. (2010). Lifting Rationality Assumptions in Binary Aggregation. *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI-2010)*.

Grandi, U., & Endriss, U. (2011). Binary Aggregation with Integrity Constraints. *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI-2011)*.

Grossi, D. (2009). Unifying Preference and Judgment Aggregation. *Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2009)*, 217–224.

Kolmogorov, A. N. (1950). *Foundations of the Theory of Probability*. Chelsea Publishing Company.

Kornhauser, L. A., & Sager, L. G. (1986). Unpacking the Court. *Yale Law Journal*, *96*, 82–117.

Larrick, R. P., & Soll, J. B. (2006). Intuitions About Combining Opinions: Misappreciation of the Averaging Principle. *Management Science*, *52*(1), 111–127.

List, C. (2003). A Possibility Theorem on Aggregation Over Multiple Interconnected Propositions. *Mathematical Social Sciences*, *45*(1), 1–13.

List, C., & Pettit, P. (2002). Aggregating Sets of Judgments: An Impossibility Result. *Economics and Philosophy*, *18*, 89–110.

List, C., & Pettit, P. (2004). Aggregating Sets of Judgments: Two Impossibility Results Compared. *Synthese*, *140*(1-2), 207–235.

Merriam-Webster. (n.d.). Congruous [Accessed: 2025-05-12].

Pauly, M., & van Hees, M. (2006). Logical Constraints on Judgement Aggregation. *Journal of Philosophical Logic*, *35*, 569–585.

Rey, S., Endriss, U., & de Haan, R. (2020). Designing Participatory Budgeting Mechanisms Grounded in Judgment Aggregation. *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning (KR)*.