Optimal Social Decision Making

Nisarg Shah

CMU-CS-16-121
August 2016

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Thesis Committee:
Ariel D. Procaccia, Chair
Nina Balcan
Avrim Blum
Tuomas Sandholm
Vincent Conitzer

Submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

This research was sponsored by the National Science Foundation under grant numbers CCF-1215883 and IIS-1350598 and Microsoft. The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of any sponsoring institution, the U.S. government or any other entity.

Keywords: Algorithmic economics, Mechanism design, Fair division, Social choice, Maximum Nash welfare, Leximin, Dynamic fair division, Implicit utilitarian voting, Robust voting

Abstract

How can computers help ordinary people make collective decisions about real-life dilemmas, like which restaurant to go to with friends, or even how to divide an inheritance? This requires the fusion of two fields: economics, which studies how people derive utility from such decisions and the incentives at play, and computer science, which informs the design of real-world computing systems that implement economic solutions. This fusion has already yielded a number of practical applications, from protection of infrastructure targets to exchange of organs.

We study two fields born of this fusion, namely *computational fair division* and *computational social choice*. Both aim to aid people in making optimal social decisions. In this thesis, we identify principled solution concepts in the two fields, and provide a comprehensive analysis of their guarantees. Our work underlies the design and implementation of the deployed service Spliddit.org, which has already been used by tens of thousands of people, and of the upcoming service RoboVote.org.

We present our results in two parts. In Part I, we focus on fair division, which addresses the question of fairly dividing a set of resources among a set of people. We study two principled solution concepts from welfare economics — maximizing Nash welfare (Chapter 2) and maximizing egalitarian welfare (a.k.a. the leximin mechanism, Chapter 3), and identify broad fair division domains in which these solutions provide compelling guarantees. Maximizing Nash welfare has been deployed on Spliddit for dividing goods (such as jewelry and art), and the leximin mechanism has applications to allocating unused classrooms in the public school system in California. We also study a specialized domain in which the leximin mechanism has made significant impact: allocation of computational resources. We build upon existing work to incorporate practicalities of real-world systems such as indivisibilities (Chapter 4) and dynamics (Chapter 5), and design mechanisms (often variants of the leximin mechanism) that provide convincing guarantees.

In Part II, we focus on social choice theory, which addresses the question of aggregating heterogeneous preferences or opinions of a set of people towards a collective outcome. We study two different paradigms of this theory: aggregation of subjective preferences towards a social consensus, and aggregation of noisy estimates of an objective ground truth towards an accurate estimate of the ground truth. In the former paradigm, we advance the recently advocated implicit utilitarian approach (Chapter 6), and offer solutions for selecting a set of outcomes; this work has been deployed on RoboVote. In the latter paradigm, we generalize the prevalent maximum likelihood estimation approach, and propose the design of *robust voting rules* that are not tailored to restrictive assumptions (Chapter 7). We also design robust voting rules for the case where the opinions of the individuals are correlated via an underlying social network (Chapter 8). Finally, taking the robustness approach to the next level, we formulate the first worst-case approach to aggregating noisy votes (Chapter 9), which has a strong connection to error-correcting codes, and show that the worst-case optimal rules offer not only attractive theoretical guarantees, but also superior performance on real data.

Acknowledgments

First, I would like to thank Ariel D. Procaccia, who, for the past five years, has not only been my advisor, but also an excellent coach and a gentle friend. You have taught me everything I know about research. When I joined CMU five years back, I was an empty book, and at each stage of my PhD career, you have inspired me with your deep insights, quick thinking, knack for finding the right research questions and directions, ability to pinpoint even the tiniest of flaws in proofs, passion for perfection, artistic writing style, incredibly original (and, apparently, quite popular!) paper titles, and the list goes on. Looking back, I still do not know how you struck the perfect balance between guiding me and allowing me think independently. Thank you for always being there when I needed you; I hope I can be at least half as good a mentor as you were.

Needless to say, this thesis would be nonexistent without the excellent people I have been fortunate to collaborate with: Yoram Bachrach, Markus Brill, Ioannis Caragiannis, Vincent Conitzer, Edith Elkind, Rupert Freeman, Albert Xin Jiang, Ian Kash, Willemien Kets, David Kurokawa, Sébastien Lahaie, Leandro Soriano Marcolino, Herv Moulin, Swaprava Nath, David C. Parkes, David M. Pennock, Ariel D. Procaccia, Yundi Qian, Sashank J. Reddi, Tuomas Sandholm, Rahul Savani, Rajiv Sethi, Eric Sodomka, Milind Tambe, Max Lee Tucker, Junxing Wang, and Yair Zick.

Special thanks to Sashank J. Reddi, for the countless hours he spent discussing my ideas, pointing out flaws in my arguments, providing fundamental insights, proofreading my papers, posing deep philosophical questions, helping me make the best career decisions, and above all, never letting me know how he managed to spare all that time, whenever I asked, from his own very successful graduate career. I also want to thank my friends Yash, Jess, Ishan, Veeru, Supraja, Shriram and Srinivas, and my peers David, Nika, Junxing, Yair, and Swaprava for their constant support, encouragement, and humorous discussions over group lunch.

Finally, to acknowledge my wife, simply for this thesis, just does not feel right, for her contributions are far greater, and to my entire life. Kinjal, thank you for making me complete on so many different levels — intellectual, philosophical, mathematical, emotional, and whatnot — in ways that I could never imagine. I dedicate, not only this thesis, but my life to you.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The alliance between computer science and economics has recently been making tremendous societal impact. Algorithms born of this synthesis, usually termed *algorithmic economics*, are already in use by major organizations for real-world applications; a few prominent examples are described below.

- *Security games* model game-theoretic allocation of security resources for protection of physical targets, used by security agencies such as the US Coast Guard and the Federal Air Marshall Service for protection of major ports and airports in the US.
- *Matching algorithms* facilitate matching of individuals or entities with each other by taking their preferences into account. Two prominent examples are *kidney exchange* algorithms, used at the United Network for Organ Sharing for conducting nationwide organ exchanges; and the *deferred acceptance* algorithm, used by the National Resident Matching Program for matching medical school students with residency programs in the US.
- *Ad auctions* facilitate efficient allocation of advertisement slots, and constitute a major revenue source for many large companies such as Google and Microsoft. *Combinatorial auctions* allow participants to bid over collections of items, and are famously used in the allocation of radio spectrum for wireless communications.

The work presented in this thesis contributes to the transition of two other subfields of algorithmic economics — computational fair division and computational social choice — from theory to practice. Informally, fair division addresses the question of fairly splitting a set of scarce resources and/or costs among a group of agents (artificial or biological); this has numerous commonplace applications such as dividing an inheritance, settling a divorce, assigning rooms among roommates and splitting the rent, etc. On the other hand, social choice theory focuses on eliciting and aggregating individual preferences or opinions towards a collective decision. It is applicable to a wide range of everyday scenarios, e.g., a group of friends selecting a restaurant or a movie to go to, the marketing team in a firm selecting a product prototype to roll out, or the citizens of a nation electing their president. At a high level, both fields address questions that arise from the social interaction among a group of people, and aim to help these people make the optimal social decisions.

Despite the enormous potential of helping millions of people every day who are already solving such problems, the real-world applications of both fields have remained severely limited. To address this, we stand on the shoulders of giants. We observe that principled economic solution concepts such as the Stackelberg equilibrium, welfare maximization, and the VCG mechanism inspired approaches that led to some of the aforementioned success stories. Such principled solution concepts have a two-fold advantage. On the one hand, they are often elegant, making it easy to convey them to the masses, and on the other hand, they are well-established across a wide range of domains, making them robust to the intricacies of a real-world setting. One wonders whether the study and use of such principled solution concepts can help advance fair division and social choice theory on the practical frontier. To that end, this thesis makes contributions on three levels.

1. We provide a comprehensive study of fundamental solution concepts from welfare economics, normative economics, and theoretical computer science in broad fair division and social choice domains, and analyze the properties of such solutions as well as the guarantees they provide.
2. When these solution concepts turn out computationally intractable ($\mathcal{NP}$-hard), we leverage various techniques from optimization, machine learning, and theoretical computer science to devise implementations for them that are sufficiently fast in practice,[1] making these solutions practicable.
3. The culmination of our work is the development of real-world systems that implement such solution concepts to help people freely access them. The next section describes, in addition to other applications, the deployment of this work to our fair division website Spliddit.org and our upcoming social choice website RoboVote.org, both of which are not-for-profit and freely accessible.

Before moving on, we remark that our work deviates from the aforementioned applications of algorithmic economics in a key aspect: while the existing applications are primarily geared towards helping large organizations make optimal decisions, our aim is to help *ordinary people* make optimal decisions in their everyday life.[2] In summary, the overarching theme of the research presented in this thesis is

> *...to develop computer programs that help groups of ordinary people make optimal social decisions.*

1.1 Real-World Impact

Spliddit.org is a not-for-profit fair division website that was launched in Nov 2014; I have been leading its development since May 2015. Spliddit's motto is to provide *provably fair solutions*. It offers (to date) five applications: dividing goods (which can be used, e.g., for inheritance division and divorce settlement), dividing chores, splitting taxi fare,

[1] We intentionally avoid the word *efficient* as it implies polynomial running time, which is neither possible nor necessary.

[2] Requests made to us for public projects, such as the project described in Chapter 3, are exceptions.

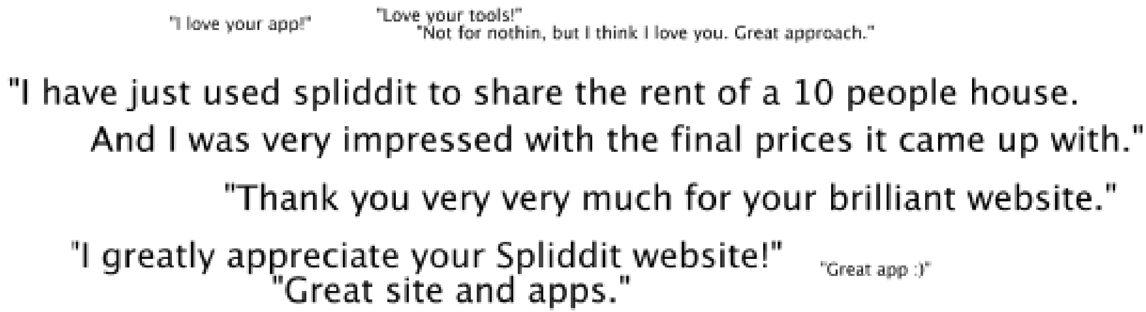

Figure 1.1: Comments from Spliddit's users

assigning rooms among roommates and splitting the rent, and assigning credit to individuals in a group project. In less than a year and a half from its launch, Spliddit has already been used by more than 70,000 people from more than 150 countries. Figure 1.1 shows a selection of comments from these users, indicative of their satisfaction.

Since the launch of Spliddit, its users have constantly been providing us useful feedback. While some users have pointed out flaws in the deployed algorithms, creating the need for us to design better solutions, other users have suggested novel real-world applications for which new solutions need to be designed. Such feedback has posed interesting and difficult new questions, and led to several serious research projects, some of which are included in this thesis.

For example, when a Spliddit user reached out to us, mentioned that the allocations returned by the goods division app did not always match his intuition of fairness, and presented an example scenario, we started searching for a better, principled solution for dividing goods. Along the way, we proved that a principled solution concept that we term the *Maximum Nash Welfare* (MNW) solution — which, in special settings, is also known under alternative names such as Competitive Equilibrium from Equal Incomes (CEEI) and proportional fairness — has compelling fairness and efficiency guarantees in a broad fair division domain, thereby generalizing classical results such as Varian's result about CEEI [195] and Weller's theorem [201]. We also devised a fast and exact implementation of the MNW solution, which has been deployed on Spliddit for allocating goods. This work is presented in Chapter 2.

Another example is the work presented in Chapter 3, which originated when a representative from the public school system in California reached out to us through Spliddit's feedback system. He asked for our help in designing a solution for the novel fair division problem of allocating unused classrooms in the public schools to the local charter schools. We observed that the literature on fair division offered another principled solution — the leximin mechanism — for several related settings [30, 44], but the exact setting of the classroom allocation problem was not previously studied. In our efforts to design a solution for the classroom allocation problem, we discovered an elegant proof for the compelling fairness, efficiency, and truthfulness guarantees of the leximin mechanism in a broad fair division domain that captures the classroom allocation setting. Our result generalizes a number of results from the literature [28, 30, 56, 99, 101, 133, 165, 186, 192, 199, 200]. We also devised an extremely fast, exact implementation for the leximin mechanism. Our work is currently under consideration for deployment by the largest school district in California.

Inspired by the success of Spliddit, since May 2015 we have also been working on the design and implementation of a new not-for-profit social choice website, RoboVote.org, which is scheduled to launch in the late summer of 2016. Like Spliddit, its aim is also to provide compelling solutions to everyday problems, albeit in the social choice realm. On a high level, RoboVote provides solutions for aggregating conflicting preferences/opinions towards a collective decision. In the objective paradigm, where a ground truth ranking of the alternatives exists (e.g., the order of different stocks by the relative change in their prices tomorrow), this means aggregation of noisy opinions towards an estimate of the ground truth. In the subjective paradigm — the classic setting where no ground truth exists, with applications to everyday scenarios such as a group of friends selecting a movie to watch or a restaurant to go to — this means aggregation of conflicting subjective preferences towards a collective choice.

The novelty of RoboVote is that it relies on optimization-based social choice approaches. For the objective paradigm, RoboVote implements voting rules that pinpoint the most likely best alternative [209], or the set most likely to contain it [179]. For the subjective paradigm, RoboVote implements the results of Boutilier et al. [35] to select a single alternative. But, previously, the extension to selecting a subset of alternatives was unavailable. Investigating this setting led us to the work described in Chapter 6, in which we established that regret minimization, which is a classic solution concept studied extensively in machine learning [27, 43], offers compelling theoretical guarantees and empirical performance on real data. We devised an implementation for this rule, and have deployed it on RoboVote.

We hope that these platforms will increase awareness of the compelling solutions that exist in both literatures for optimally solving everyday problems, enable their widespread use, and in time fuel and guide further research in both fields.

1.2 Background

Before diving into the technical details, let us present a brief overview of the basic concepts and related lines of work in fair division and social choice. On a high level, both fields aim to find "socially desirable outcomes" in settings that involve an interaction between multiple agents. But, what does socially desirable mean? The answer to that depends on the specific setting of interest, as we describe below.

1.2.1 Fair Division

The basic fair division question is: *How do we fairly divide a set of resources $\mathcal{R}$ among a set of people $\mathcal{N}$?* Following the traditional terminology, we will henceforth use the terms *players* and *goods*[3] to refer to people (or other types of agents) and resources, respectively.

Types of goods: In general, $\mathcal{R}$ can be a finite, countably infinite, or even uncountably infinite set. This can model a mixture of divisible goods (which can be split among players) and indivisible goods (which must be allocated entirely to a single player).

Allocation: An allocation (a.k.a. division, assignment) $\boldsymbol{A} = (A_i)_{i \in \mathcal{N}}$ maps each player $i \in \mathcal{N}$ to a (mutually exclusive) subset of goods $A_i \in 2^{\mathcal{R}}$, which is called the player's *bundle*. Certain settings may impose additional restrictions on feasible allocations. For instance, in fairly scheduling presentations at a conference, it is required that each presenter be assigned exactly one time slot. Let $\mathcal{A}$ denote the set of feasible allocations.

Preferences/valuations: To measure how good an allocation is, we need information about how much players like their bundles. A common approach is to use *cardinal* preferences: the valuation (a.k.a. utility function) of player i, denoted $v_i : 2^{\mathcal{R}} \to \mathbb{R}$, describes the real-valued utility derived by the player for each possible bundle of goods.[4] Pareto [162] advocated the use of *ordinal* preferences, which compare bundles instead of assigning them a real-valued utility.[5]

There is an obvious tradeoff between the use of ordinal versus cardinal preferences. Ordinal preferences reveal less information than cardinal preferences, but also impose less cognitive load on the players. The use of ordinal preferences is more prevalent in fields such as social choice where it is difficult to assign real values to outcomes (e.g., to the event that one's preferred candidate becomes the president). In contrast, in many practical applications of fair division, people can determine the amount of money they are willing to pay to receive a bundle of goods (e.g., a chair and a desk), and can submit that as the value. Hence, the fair division literature often uses cardinal preferences.

That said, much of the literature intentionally avoids interpersonal comparisons between the valuations of two players. Moulin [151, pp. 6-7] provides an excellent account of the technical as well as moral objections against interpersonal comparisons. In fact, several fairness and efficiency desiderata only compare bundles from the perspective of the same player. This makes the assumption of having cardinal valuations unrestrictive as the only information really used is ordinal information.

Mechanism: A mechanism f takes as input the reported valuations (a.k.a. revelations) of the players $\boldsymbol{v} = (v_i)_{i \in \mathcal{N}}$, and returns an allocation $\boldsymbol{A} = (A_i)_{i \in \mathcal{N}}$. This process can

[3] We require them to be "goods", i.e., be valued positively by the players. Fair allocation of "bads" (i.e., items with negative utility) is the subject of exciting ongoing research.

[4] Typically, we assume $v_i(\emptyset) = 0$ and $v_i(X) \leqslant v_i(Y)$ for all $X \subseteq Y \subseteq \mathcal{R}$.

[5] Restrictions on feasible allocations can often be modeled using the valuations of the players. For instance, in the previous example of scheduling conference presentations, we can modify players' valuations such that they do not have additional value for receiving multiple time slots. Together with the mild restriction (called *nonwastefulness*) that players should not be allocated extraneous resources, this reconstructs the desired restriction.

be broken into three stages: eliciting the valuations, setting the fairness and efficiency desiderata, and finding an allocation that satisfies these desiderata.

Step 1: Eliciting the valuations.

When players are self interested, this can be tricky as a player may misrepresent her valuation to receive a more preferred bundle. Typically, monetary transfers from the designer to the players are sufficient to incentivize the players to report truthfully [58, 106, 196]. However, numerous settings prohibit the use of money by normative considerations (see, e.g., the classroom allocation application in Chapter 3). In such cases, incentivizing truth-telling without the use of money is an often sought-after property, which is defined below.

Definition 1.1 ((Group) Strategyproofness)**.** A mechanism is called strategyproof if no player can gain by lying about her valuation, i.e., if truth-telling is a weakly dominant strategy for every player. In other words, for each player $i \in \mathcal{N}$, if $\boldsymbol{A}$ (resp. $\boldsymbol{A'}$) is the allocation returned when player i reports her true valuation v_i (resp. a false valuation v_i'), then, ceteris paribus, we must have $v_i(A_i) \geqslant v_i(A_i')$. A mechanism is called group strategyproof if no group of players can simultaneously gain by misreport their valuations. A mechanism is called strongly group strategyproof if no group of players can simultaneously misreport their valuations in a way such that no group member is worse off, and at least one group member is strictly better off.

Step 2: Defining fairness and efficiency desiderata.

The next step is to define what "socially desirable" means in the fair division context. Vaguely, we want our allocation to be both fair and efficient. The standard approach, derived from normative economics, proposes a scale of fairness and efficiency desiderata, which are often defined in a general form, and are applicable to a wide range of settings.

The most compelling of all fairness notions — perhaps the gold standard — is called envy freeness, and was originally proposed by Foley [93].

Definition 1.2 (Envy-freeness)**.** An allocation $\boldsymbol{A} \in \mathcal{A}$ is called envy free if every player values her bundle at least as much as every other player's bundle, that is, if

$$v_i(A_i) \geqslant v_i(A_j),\ \forall i, j \in \mathcal{N}.$$

In some sense, envy-freeness requires everyone to be simultaneously be happy about the allocation, which is not always possible. For instance, if an indivisible diamond needs to be allocated between two players, the one receiving it will inevitably be envied by the other. However, when envy-freeness can be achieved, it is often sufficient, and implies many other fairness notions. Let us review another popular notion of fairness.

Definition 1.3 (Proportionality)**.** An allocation $\boldsymbol{A} \in \mathcal{A}$ is called proportional if every player's utility is at least $1/|\mathcal{N}|$ fraction of her utility for the set of all resources $\mathcal{R}$, that is, if

$$v_i(A_i) \geqslant \frac{1}{|\mathcal{N}|} \cdot v_i(\mathcal{R}),\ \forall i \in \mathcal{N}.$$

Proportionality declares $(1/|\mathcal{N}|) \cdot v_i(\mathcal{R})$ as the fair share of player i,[6] and demands each player receive at least her fair share. Observe that in the diamond-allocation example above, proportionality is also impossible to achieve. Assuming all resources are allocated, it can be shown that envy-freeness implies proportionality for additive valuations.

Budish [44] (building on concepts introduced by Moulin [150]) presented an alternative fairness notion. It is based on the classic cut-and-choose procedure for two players, in which a player divides the resources into two bundles, but gets to pick her bundle after the other player. With divisible goods, this ensures envy-freeness. Budish generalized this to the case of more than two players and possible indivisibilities.

Definition 1.4 (Maximin Share)**.** Let $\Pi_k(\mathcal{R})$ denote the labeled set of partitions of $\mathcal{R}$ into k bundles. The maximin share (MMS) guarantee of player $i \in \mathcal{N}$, denoted $\mathtt{MMS}_i$, is the utility the player can guarantee herself if she divides the set of resources into n bundles, but picks her bundle last and receives the worst one according to her:

$$\mathtt{MMS}_i = \max_{\boldsymbol{A} \in \Pi_{|\mathcal{N}|}(\mathcal{R})} \min_{k \in [|\mathcal{N}|]} v_i(A_k).$$

An allocation $\boldsymbol{A} \in \mathcal{A}$ is called an α-MMS allocation if

$$v_i(A_i) \geqslant \alpha \cdot \mathtt{MMS}_i, \ \forall i \in \mathcal{N}.$$

One can check that for additive valuations, maximin share is implied by proportionality. Let us now review another relaxation of envy-freeness, which has been studied in different forms over the years [44, 50, 105, 165].

Definition 1.5 (Envy-freeness up to one good)**.** An allocation $\boldsymbol{A} \in \mathcal{A}$ is called envy free up to one good if every player values her bundle at least as much as any other player's bundle after removing at most one good from the other player's bundle, that is, if

$$\Big(v_i(A_i) \geqslant v_i(A_j)\Big) \vee \Big(\exists r \in A_j,\ v_i(A_i) \geqslant v_i(A_j \setminus \{r\}\Big),\ \forall i, j \in \mathcal{N}.$$

In addition to fairness, we often desire that resources should not be wasted, i.e., the allocation should be efficient. Below, we review a popular notion of efficiency.

Definition 1.6 (Pareto Optimality)**.** An allocation $\boldsymbol{A} \in \mathcal{A}$ is called Pareto optimal if no alternative allocation can increase the utility to a player without reducing the utility to another player, that is, if

$$\Big(\exists i \in \mathcal{N},\ v_i(A'_i) > v_i(A_i)\Big) \implies \Big(\exists j \in \mathcal{N},\ v_j(A'_j) < v_j(A_j)\Big),\ \forall \boldsymbol{A'} \in \mathcal{A}.$$

Step 3: Computing a fair and efficient allocation.

We now review standard fair division settings, and compelling mechanisms for them.

[6]This would be her utility if $\mathcal{R}$ is given entirely to a player chosen uniformly at random.

Additive valuations: Perhaps the most studied class of valuations is additive valuations, in which each player's value for a bundle of goods is the sum[7] of her values for the individual goods in the bundle. Let us review two special cases of this setting.

Divisible goods: If the goods are divisible (i.e., can be split among players), we can simply model a single, heterogeneous[8] good as one can "embed" all goods in it. Such a good is traditionally called a *cake*. Cake cutting studies fair and efficient allocations of a cake; see the excellent survey and book chapter by Procaccia [169, 170] for a detailed review.

There are two existing results that are particularly relevant for us. Settling a long standing open question, Aziz and Mackenzie [11] recently presented a bounded time protocol to compute an envy-free division of the cake (that allocates the entire cake). However, if we were to relax the computational restriction, Weller [201] had already shown that a principled solution concept called *Competitive Equilibrium from Equal Incomes* (CEEI) [98, 113, 195] finds an allocation that is both envy free and Pareto optimal.

Indivisible goods: Let us now consider the setting where all the goods are indivisible. Interestingly, barring computational considerations, this is a strictly more general setting because partitioning a divisible good into k identical indivisible parts, and taking the limit of k going to infinity can model the divisible good perfectly. Unfortunately, with indivisibilities, a CEEI allocation may not exist. Hylland and Zeckhauser [113] showed that CEEI can still be used to compute a lottery for allocating n players to an equal number of indivisible "slots" in an envy free and ex-ante Pareto optimal way. Budish [44] proposed an approximate CEEI that is envy free up to one good and approximately Pareto optimal.[9] A few years later, settling an open question, Procaccia and Wang [175] showed that, while it is *not* always possible to find an MMS allocation, one can always find a $(2/3)$-MMS allocation. This was the first approach that provided a non-trivial, provable fairness guarantee for allocating indivisible goods.

Two years later, Caragiannis et al. [50] showed that a much simpler and principled mechanism, *Maximum Nash Welfare*, that maximizes the *product* of the players' utilities satisfies a more elegant relaxation of envy-freeness — envy-freeness up to one good — and Pareto optimality. Given that maximum Nash welfare coincides with CEEI for divisible goods, it follows that their result generalizes results by Varian [195] and Weller [201]. Further, maximum Nash welfare is also shown to be $(1/\Theta(\sqrt{|\mathcal{N}|}))$-MMS.

Non-additive valuations: For dichotomous preferences, under which players have values 0 or 1 for each feasible allocation, Bogomolnaia and Moulin [30] showed that another principled mechanism called *the leximin mechanism* has a number of compelling properties: proportionality, envy freeness, Pareto optimality, and group strategyproofness. Ghodsi et al. [99] studied the leximin mechanism for allocating divisible computational resources in a cluster, under the name of *Dominant Resource Fairness* mechanism, and showed the exact same properties. Later, Kurokawa et al. [126] generalized these

[7]For divisible goods, sum becomes integration.

[8]"Heterogeneous" means different parts of the good may be valued differently by the players.

[9]In addition, it satisfies a weaker version of the maximin share guarantee.

as well as several other results from the literature, showing that the properties of the leximin mechanism hold for a much more general class of valuations.

Principled economic solution concepts for fair division.

At its core, fair division aids social decision making, and the goal is to improve overall well-being of the group of individuals involved. While the individual well-being is traditionally defined using a utility function, there is no objective definition of overall well-being of a group. To that end, a fundamental approach from welfare economics suggests choosing the decision that optimizes a *social welfare function*, a monotonic function of the individual utilities that serves as a measure of overall well-being of the group. Three principled mechanisms that maximize well-established social welfare functions are:

- *maximizing the utilitarian welfare* [31, 94], which measures the *sum* of the utilities to the individuals;
- *maximizing the Nash welfare* [50, 59, 98, 113, 130, 195, 201], which measures the *product* of the utilities to the individuals; and
- *maximizing the egalitarian welfare*, which measures the *minimum* of the utilities to the individuals. The well-known *leximin mechanism* [28, 30, 56, 99, 101, 126, 165, 186] is a refinement that not only maximizes the minimum utility, but also breaks ties in favor of allocations with higher second minimum utility, breaks remaining ties in favor of allocations with higher third minimum utility, and so on.

1.2.2 Social Choice

The basic social choice question is the following: *How do select, in a socially desirable fashion, an outcome from a set of possible outcomes given heterogeneous preferences/opinions of a set of people over the outcomes?* Following the traditional notation from social choice theory, we will use the term *voters* to denote people, *votes* to denote their preferences/opinions, *profile* to denote the collection of votes, and *voting rule* to denote the mapping from profiles to selected outcomes. As discussed in Section 1.2.1, in many social choice settings, it is difficult for voters to assign real-valued utilities to alternatives. Hence, the literature often revolves around ordinal. Unless mentioned otherwise, we will henceforth assume that each voter provides a *ranking* (a.k.a. total order) over the set of alternatives. Before diving into approaches to designing voting rules, let us dip our nib into the history of social choice theory.

A brief history of social choice theory: The origin of the field dates back to the late eighteenth century, when Condorcet [60] proposed the *majority rule* to aggregate ranked preferences into a collective decision. According to his proposal, the decision between each pair of alternatives should be made by considering which alternative a majority of the voters prefer. The collective decision should then be made by following such pairwise decisions. Unfortunately, a classic paradox shows that the will of the majority on pairwise decisions can result in cycles.[10]

[10]That is, it may happen that for alternatives a, b, and c, a majority prefers a to b, and b to c, but c to a.

Approximately two centuries later, Arrow [5] gave birth to modern social choice theory with his classic impossibility result. His approach was inspired from normative economics. He formulated three axioms that "socially desirable" voting rules should satisfy, and showed no voting rule satisfies them together. Subsequent research relaxed the axioms to obtain positive results. An important step in the early days of modern social choice theory was formalization of Condorcet's ideas. One formalization was in the form of Kemeny's rule [119, 209], which finds minimal changes in the will of the majority to eliminate such cycles. Another formalization was Dodgson's rule [77], which finds minimal changes in the input preferences until one alternative is preferred to every other alternative by a majority.

The next breakthrough was the birth of computational social choice due to papers by Bartholdi et al. [19, 20]. They infused the notion of computational hardness from theoretical computer science into social choice theory, and showed that while it can be problematic, as determining the winner may sometimes be hard, it can also be a boon, as one can prevent voters from efficiently manipulating voting rules, thus circumventing another classic impossibility result by Gibbard [102] and Satterthwaite [189]. Subsequently, research on computational social choice boomed with numerous results on complexity or computability of winner determination [22, 68, 79, 142], strategic voting [18, 64, 67, 69, 76, 177], bribery and control [21, 88, 89, 90, 112], etc.

An alternative paradigm and ground truth: While traditionally social choice theory addresses aggregation of subjective preferences, a new paradigm has come to light with the recent advent of crowdsourcing platforms. In this alternative paradigm, whose roots date back to Condorcet [60], there exists an unknown, objective ground truth comparison among the alternatives. Voters submit their noisy opinions about this ground truth, and the goal now is to aggregate these noisy opinions to uncover the ground truth, or the best estimate thereof.

Although this is a conceptually different question, which does not require a "socially desirable" aggregation, at its core the technical problem is still the same: aggregating rankings of the alternatives to find a central ranking, of sorts. In fact, applications in the subjective and objective domains differ subtly; e.g., in case of multiple candidates running for president, asking voters to rank candidates according to their preference falls within the subjective domain, whereas asking voters to rank the candidates by their likelihood of becoming the president falls within the objective domain.

Approaches to designing voting rules: Currently, there are four prominent approaches to designing voting rules.

1. *The axiomatic approach.* In this approach, one defines axiomatic desiderata that must always be satisfied during the aggregation. Examples of such desiderata include majority consistency, Condorcet consistency, monotonicity, independence of irrelevant alternatives, and consistency. Refer to the book chapters by Brandt et al. [39] and Zwicker [213] for a detailed exposition.
2. *The distance rationalizability approach.* Instead of defining desiderata that must always be satisfied, this approach identifies a subset of profiles in which an obvious

collective choice exists. Then, given an arbitrary profile, one finds the nearest profile (in a precise sense) in which an obvious collective choice exists, and returns this choice. Refer to the book chapter by Elkind and Slinko [83] for an overview.

3. *The implicit utilitarianism approach.* A recent line of papers [35, 51, 171] have put forward a fundamentally different approach. They assume that the expressed ranked preferences are proxy for underlying cardinal utilities, and attempt to maximize the (cardinal) utilitarian welfare. Because it is impossible to choose the welfare-maximizing outcome given the limited ordinal information, they choose the outcome that provides the best worst-case approximation of welfare.
4. *The statistical framework.* Inspired by the objective domain, the statistical framework for voting models the generation of noisy opinions given the ground truth, and uses the observed opinions to predict what the ground truth must have been. The standard approach in this framework is the maximum likelihood estimation (MLE) approach [66, 82, 85, 140, 172, 179], which uses a single model of noisy opinion generation to identify and return the collective choice that is *most likely* to have generated the observed opinions. Recently, Caragiannis et al. [48, 49] extended this approach to designing *robust voting rules*, which identify the collective choice that is likely to have generated the observed opinions according to a wide family of generation models rather than being tailored to a single, often inaccurate generation model. Procaccia et al. [182] proposed the first worst-case approach to voting, which is more robust and fundamentally different than the statistical framework.

In summary, social choice theory has a somewhat different take on what "socially desirable" means, but, unlike fair division, several fundamentally different answers exist: accuracy in pinpointing the ground truth, satisfaction of axiomatic desiderata, welfare maximization, etc. We note that computational aspects (e.g., of winner determination) are important and have been studied in all of the abovementioned approaches.

1.3 Overview of Thesis Structure and Our Results

The thesis presents a selection of my work on fair division and social choice theory. To keep the thesis succinct, only works that have novel conceptual contributions, significantly advance the understanding of an existing concept, and/or have important real-world implications are included. The content is organized into two parts, each consisting of four chapters: Part I describes the work on fair division, and Part II describes the work on social choice theory.

Part I: Fair Division

In this part, we present our analysis of two fundamental solution concepts — maximizing the Nash welfare and maximizing the egalitarian welfare (more specifically, the leximin mechanism), as well as of solution concepts derived from them.

Chapter 2: Allocating Goods and Maximum Nash Welfare

This chapter invokes the notion of envy-freeness up to one good (EF1) for allocation of indivisible goods under additive valuations. This notion was first introduced by Lipton et al. [134], albeit in a more general context. The algorithm they presented satisfies EF1, but does not guarantee additional properties such as Pareto optimality (PO). Budish [44] presented approximate CEEI that satisfies EF1, but is only approximately efficient. Further, the guarantees provided by his mechanism are not compelling when the number of players is small and there is a single copy of each good, which is typical on Spliddit.

The solution concept of our interest, maximizing Nash welfare, is a well-established solution concept. It was first studied by Nash [154] in the context of his classic bargaining problem. In the networking community, the same solution goes by the name of *proportional fairness*, due to another property that it satisfies when goods are divisible [118]. In the context of allocating indivisible goods, Ramezani and Endriss [184] showed that the problem is $\mathcal{NP}$-hard for additive valuations. Cole and Gkatzelis [59] gave a constant-factor, polynomial-time approximation algorithm, but such an approximation need not satisfy any of the fairness or efficiency guarantees introduced in Section 1.2.1. The APX-hardness result by Lee [130] shows that a constant-factor approximation is the best one can hope for if we impose the restriction of polynomial running time.

In this chapter, we establish fairness and efficiency properties of the Maximum Nash Welfare (MNW) solution for allocating indivisible goods (which, as we saw in Section 1.2.1, is strictly more general than allocating divisible goods). Our main result is that it always outputs an allocation that is both EF1 and PO. And while EF1 and PO are straightforward to obtain in isolation, the MNW solution is the only known mechanism in the literature to achieve them together, which is a strong argument in its favor. Further, guarantees such as EF1 and PO are easy to convey to the end users, which is crucial for the deployment of the solution concept to Spliddit.

Additionally, we also show that the MNW solution gives a tight $2/(1+\sqrt{4n-3})$ approximation to the maximin share guarantee, and a tight $(\sqrt{5}-1)/2 \approx 0.618$ approximation to another fairness desideratum that we term the pairwise maximin share guarantee, where n is the number of players. The approximation ratios are, however, 1 in more than 90% of the instances in our experiments on real data from Spliddit. We also show that the fairness and efficiency guarantees of the MNW solution extend, to some extent, to the more general class of submodular valuations.

Our final contribution is a fast implementation for exactly computing the MNW solution for the form of valuations elicited on Spliddit, in which a player is required to divide 1000 points among the available goods.[11] Our algorithm scales very well, solving relatively large instances with 50 players and 150 goods in less than 30 seconds, while other candidate algorithms we describe fail to solve even small instances with 5 players and 15 goods in twice as much time.

Chapter 3: Allocating to Strategic Agents and The Leximin Mechanism

[11] This is still an $\mathcal{NP}$-hard problem due to a result by Nguyen et al. [155].

In this chapter, we study a real-world fair division setting in which unused classrooms in the public schools in California are to be allocated fairly to the local charter schools. More specifically, we have i) public schools (a.k.a. facilities), each of which has a given number of unused classrooms — its *capacity*, and ii) charter schools (a.k.a. agents), each of which requires a certain number of classrooms — its *demand*, views a subset of facilities as acceptable, and has dichotomous preferences, i.e., has utility is 1 if it receives exactly the desired number of classrooms at a single acceptable facility and 0 otherwise.

Previous work fails to provide compelling guarantees in this setting. The pseudo-market mechanism by Hylland and Zeckhauser [113] and the probabilistic serial mechanism by Bogomolnaia and Moulin [29] are not strategyproof, and only work when the number of agents and goods (i.e., classrooms) are equal. The probabilistic serial mechanism admits several extensions for multi-unit demands [10, 45, 55, 123, 183], but they can only allocate *at most* d classrooms to a charter school demanding d classrooms, whereas we require the charter school to receive either exactly d classrooms or no classrooms at all. Approximate CEEI due to Budish [44] has practical guarantees, but only if each facility has a large number of unused classrooms, which is not the case in practice. The mechanism that uniformly randomizes among allocations maximizing the utilitarian welfare is known to be envy free, Pareto optimal, and strategyproof in a setting more general than the classroom allocation setting [31, 94], but violates proportionality and suffers from "tyranny of the majority", making it highly undesirable in our setting; refer to Chapter 3 for more details.

The closest starting point is the leximin mechanism, which Bogomolnaia and Moulin [30] studied for a special case of our setting, in which all demands and capacities are 1 and the number of agents and facilities are equal, and showed that it satisfies proportionality, envy-freeness, Pareto optimality, and group strategyproofness. Their techniques, however, do not extend directly to our more general setting.

In Chapter 3, we show that the leximin mechanism remains compelling (i.e., satisfies all the four properties) in the classroom allocation setting. The beauty of these properties, as well as of the leximin mechanism itself, is that they are intuitive and can easily be explained to a layperson; this contributes towards the practicability of the approach. Importantly, we establish the properties of the leximin mechanism in a setting much more general than the classroom allocation setting, thereby generalizing results from a vast number of papers in the literature [28, 30, 56, 99, 101, 165, 186, 192, 199, 200] on topics including resource allocation, cake cutting, and kidney exchange.

Studying the mechanism from a combinatorial optimization viewpoint, we show that the expected number of classrooms allocated by the leximin mechanism is always at least $1/4$ of the maximum number of classrooms that can possibly be allocated. We conjecture that an improved bound of $1/2$ is feasible.

Our final contribution is an extremely fast implementation for the $\mathcal{NP}$-hard problem of exactly computing the leximin allocation. The crux of our implementation is to modify a naïve approach, which solves a sequence of linear programs with an exponential number of variables, work with the duals that have an exponential number of constraints, and formulate a separation oracle for them as an integer linear program. In our experiments based on real data, our implementation solves instances with 300 charter schools and more than 1000 public schools (which is larger than any real-world instance) in a few minutes on average. Remarkably, in the experiments, the expected number of charter schools satisfied and the expected number of classrooms allocated by the leximin mechanism are, on average, at least 98% of the respective maxima.

Chapter 4: Fair Division of Computational Resources

A prominent application of the leximin mechanism studied in Chapter 3 is for allocating divisible computational resources, e.g., CPU, RAM, and network bandwidth to agents (such as jobs and users) in computing systems (such as operating systems and computational clusters). The traditional approach employed a *single resource* abstraction, in which fixed amounts of different resources were bundled into "slots", and the slots are then allocated to the agents. However, in a realistic environment where agents have heterogeneous demands, such an abstraction inevitably leads to significant inefficiencies.

Ghodsi et al. [99] suggested modeling the heterogeneous demands of the agents through Leontief preferences that reflect the desire to maintain a fixed proportion between resource types. For example, if an agent wishes to execute multiple instances of a job that requires 2 CPUs and 1 GB RAM, it would prefer 3 CPUs and 1.5 GB RAM over 2 CPUs and 1 GB RAM, but would be indifferent between the former allocation and 4 CPUs and 1.5 GB RAM, because in both instances, the agent can only run 1.5 instances of its task (allowing divisible tasks). They studied the leximin mechanism, under the name of the Dominant Resource Fairness (DRF) mechanism, and showed that for Leontief preferences, it satisfies four compelling desiderata: sharing incentives (SI, a.k.a. proportionality), envy freeness (EF), Pareto optimality (PO), and strategyproofness (SP). However, their work suffers from two unrealistic assumptions — divisible tasks and strictly positive demands (i.e., each agent requires a positive quantity of each resource).

In this chapter, we make three key contributions. First, we consider a realistic extension of the Ghodsi et al. [99] setting that incorporates priorities among agents in the form of agent weights, and relaxes their restrictive assumption that every agent requires each resource. We show that all four properties of the DRF mechanism hold (and strategyproofness can in fact be strengthened to group strategyproofness) in this more general setting. We note that our later work [126] presented in Chapter 3 subsumes this analysis as part of a more general analysis.

Second, we explore the relation between the aforementioned desirable properties, and maximization of the social welfare—the sum of utilities of the agents. To study social welfare we must assume an *interpersonal comparison of utilities*. Focusing on a natural utility function, we observe that DRF can produce allocations that only provide roughly $1/m$ of the social welfare of the optimal allocation, where m is the number of resources. However, we vindicate DRF, demonstrating that this poor welfare property is necessary for any mechanism that satisfies *at least one* of the three properties SI, EF, and SP.

Finally, we tackle another realistic extension in which the tasks in the example presented above are indivisible, that is, an agent's task would require a minimum, *indivisible* bundle of resources. For example, if an agent requires 2 CPUs and 1 GB RAM to run one instance of its task, allocating 1 CPU and 1/2 GB RAM would be no more preferred than allocating nothing at all. In this novel setting, we observe that DRF performs poorly; in fact, envy freeness and Pareto optimality are trivially incompatible. Inspired by the aforementioned envy-freeness up to one good property, we introduce a similar property *envy-freeness up to one bundle (EF1)*: an agent i prefers its own allocation to the allocation of another agent j, given that a single copy of the demanded bundle of i is removed from the allocation of j. While SP is still incompatible with PO and EF1 (as well as with PO and SI) in this setting, we design a mechanism, SEQUENTIALMINMAX, which satisfies PO, SI, and EF1. The mechanism is a variant of the leximin mechanism that sequentially allocates bundles to minimize the maximum share of a resource that any agent has after allocation.

In related work, Li and Xue [131] provided a characterization of mechanisms that satisfy certain desirable properties under Leontief preferences, but their results do not capture DRF and also suffer from the strictly positive demands assumption. Friedman et al. [97] introduced a family of *weighted* DRF mechanisms, but the weights in their model are internal to the mechanism, and do not reflect agent priorities. Further, they, too, assume strictly positive demands. Dolev et al. [78] studied an alternative fairness criterion, *no justified complaints*, showed existence of allocations satisfying it, and compared them with the DRF allocation. Gutman and Nisan [108] presented polynomial-time algorithms for computing allocations under a family of mechanisms that includes DRF, and for computing allocations satisfying no justified complaints.

Chapter 5: Dynamic Fair Division

A major limitation of the existing work on fairly dividing computational resources, and in fact, of the entire existing literature on fair division, is that they study one-shot fair allocations in static settings where all the agents and all the resources exist in the system from the beginning. Most real systems are more dynamic: agents typically arrive and depart over time, and the system adjusts the allocation of resources accordingly.

The networking community has studied the related problem of fairly allocating a single homogeneous resource in a queuing model where each agent's task requires a given number of time units to be processed. In other words, in these models tasks are processed over time, but demands stay fixed, and there are no other dynamics such as agent arrivals and departures. The well-known *fair queuing* solution [73] allocates one unit per agent in successive round-robin fashion. This solution has also been analyzed by economists [152].

In the fair division literature, the only previous attempt to study fairness in a dynamic environment was by Walsh [198], who proposed the problem of fair online cake cutting where agents arrive, take a piece of cake, and immediately depart. He suggested several desirable properties for cake cutting mechanisms in this setting, and showed that adaptations of classic mechanisms achieve these properties. However, as Walsh pointed out, allocating the whole cake to the first agent also achieves the same properties, making the setting trivial. His notion of *forward envy freeness* is related to our notion of *dynamic envy freeness*.

In this paper, we extend the model introduced in Chapter 4 to include agent arrivals over time, thus introducing the first non-trivial model for and initiating the field of dynamic fair division.

Even on the conceptual level, dynamic settings challenge some of the premises of fair division theory. For example, if one agent arrives before another, the first agent should intuitively have priority; what does fairness mean in this context? We introduce the concepts that are necessary to answer this question, and design novel mechanisms that satisfy our proposed desiderata.

Specifically, we consider an environment in which agents arrive over time (but do not depart), and the allocations made to agents are irrevocable, i.e., the mechanism can allocate more resources to an agent over time, but cannot take resources back. See Chapter 5 for additional discussion on these assumptions.

We adapt prominent notions of fairness, efficiency, and truthfulness to our dynamic settings. For fairness, we ask for envy freeness (EF) and sharing incentives (SI) at each point in time. For truthfulness, we seek strategyproof (SP) mechanisms under which agents cannot gain (at any point in time) from misreporting their demands. For efficiency, we introduce the notion of *dynamic Pareto optimality* (DPO): if k agents are entitled to k/n of each resource, the allocation should not be dominated (in a sense that will be formalized in the chapter) by allocations that divide these entitlements.

Our first result is an impossibility: DPO and EF are incompatible. We proceed by relaxing each of these properties. First, we relax EF to a new dynamic property, which we call *dynamic EF (DEF)*, that allows an agent to envy another agent that arrived earlier, as long as the former agent was not allocated resources after the latter agent's arrival. We construct a new mechanism, DYNAMIC DRF, and prove that it satisfies SI, DEF, SP, and DPO.

Next, we relax the DPO property. Our *cautious DPO* (CDPO) notion allows allocations to only compete with allocations that can ultimately guarantee EF, regardless of the demands of future agents. We design a mechanism called CAUTIOUS LP, and show that it satisfies SI, EF, SP, and CDPO. In a sense, our theoretical results are tight: EF and DPO are incompatible, but relaxing only one of these two properties is sufficient to enable mechanisms that satisfy both, in conjunction with SI and SP.

Despite the assumptions imposed by our theoretical model, we believe that our new mechanisms are compelling, useful guides for the design of practical resource allocation mechanisms in realistic settings. Indeed, we test our mechanisms on real data obtained from a trace of workloads on a Google cluster, and obtain encouraging results.

Part II: Social Choice Theory

In this part, we study voting rules in both the subjective and the objective domains. Chapter 6 studies a rather modern, implicit utilitarianism approach for the subjective domain. Chapters 7 and 8 introduce the concept of robust voting rules, and design them. Chapter 9 takes the robustness approach further by introducing a worst-case model of voting.

Chapter 6: Subset Selection Using Implicit Utilitarian Voting

In this chapter, we study the classic social choice problem of aggregating ranked preferences of a set of voters into a collective decision. While traditional social choice theory typically takes a normative approach, focusing on the design of voting rules that satisfy certain desirable axioms, researchers in computational social choice [40] often advocate *quantitative* approaches. The high-level idea is to identify a compelling objective function, and design voting rules that optimize this function.

A recent line of papers [3, 4, 35, 47, 171] advocated an approach in which one assumes that voters have latent cardinal utilities which are consistent with the expressed ranked preferences, and focuses on maximizing the *utilitarian social welfare*, i.e., the sum of voters' utilities. We refer to this approach as *implicit utilitarian voting*. In this approach, the performance of a voting rule can be quantified via a measure called *distortion*: the worst-case (over utility functions consistent with the reported profile of rankings) ratio between the social welfare of the optimal (welfare-maximizing) alternative, and the social welfare of the alternative selected by the voting rule.

Procaccia and Rosenschein [171] analyzed the distortion of existing voting rules, and Boutilier et al. [35] designed voting rules that minimize distortion. The work of Boutilier et al. [35] provides a compelling solution but only when a single alternative is selected by the voting rule. Several common applications (e.g., choosing a committee) require selection of a subset of alternatives. In this paper, we extend the implicit utilitarian approach to selecting a subset of alternatives, and understand the guarantees they provide, as well as their performance in practice.

There exist other approaches in the computational social choice literature for subset selection. Under the Chamberlin-Courant method, for example, each voter assigns a score to a set equal to the highest score of any alternative in the set, and the (computationally hard) objective is to choose a subset of size k that maximizes the sum of scores [54, 178]. Skowron et al. [191] generalize the way in which the score of a voter for a subset of alternatives is computed. Aziz et al. [12] propose selecting a subset of alternatives in order to satisfy a fairness axiom they term *justified representation*, and study whether common voting rules satisfy this axiom. The budgeted social choice framework of Lu and Boutilier [138] is more general in that the number of alternatives to be selected is not fixed; rather, each alternative has a cost that must be paid to add it to the selection.

We make four main contributions. First, on a conceptual level, we introduce the additive notion of *regret* into the implicit utilitarian voting setting, as an alternative to the multiplicative notion of distortion. Second, we derive worst-case bounds on the distortion and regret of optimal deterministic and randomized voting rules. Third, we empirically compare the *worst-case-optimal* deterministic voting rules with respect to distortion and regret — denoted $f^*_{\texttt{dist}}$ and $f^*_{\texttt{reg}}$, respectively — with a slew of well-known voting rules, in terms of the *average-case* distortion and regret using synthetic and real data. We find that $f^*_{\texttt{reg}}$ outperforms all other rules on average, even when measuring distortion! Fourth, we develop a scalable implementation for $f^*_{\texttt{reg}}$ (which, we show, is $\mathcal{NP}$-hard to compute). We have deployed this rule to our upcoming social choice website RoboVote.org.

Chapter 7: Robust Voting Rules

In this chapter, we study the design of voting rules in the objective domain of social choice, i.e., in the case where a latent ground truth exists. Specifically, we assume that there exists a true ranking of the alternatives, and we receive ranked opinions which are noisy estimates of this underlying true ranking. Our goal is to recover the true ranking.

The standard approach in the literature is the *maximum likelihood estimation* (MLE) approach [7, 9, 65, 66, 70, 82, 140, 172, 179, 206, 207, 209]. In this approach, one fixes a "noise model", which defines the probability of each ranking being generated as a noisy estimate given each possible ground truth ranking, and uses this noise model to identify the ranking that has the most likelihood of being the ground truth. A classic noise model is Mallows' model [60, 139], under which each voter ranks each pair of alternatives correctly with probability $p > 1/2$ and incorrectly with probability $1 - p$, and the mistakes are i.i.d.[12] The MLE rule for Mallows' model is known as the Kemeny rule [119].

As compelling as the MLE approach is, there are many different considerations in choosing a voting rule, and insisting that the voting rule be an MLE is a tall order (there is only one MLE per noise model); this is reflected in existing negative results [66, 85]. We relax this requirement by asking: *How many* votes do prominent voting rules need to recover the true ranking with high probability? In crowdsourcing tasks, for example, the required number of votes directly translates to the amount of time and money one must spend to obtain accurate results. Clearly, the MLE rule requires the least number of votes. Focusing on Mallows' model, we show that its MLE rule (the Kemeny rule) sits within a larger family of voting rules — pairwise-majority consistent (PM-c) rules — that all require asymptotically optimal number of votes (logarithmic in the number of alternatives) to pinpoint the ground truth with a desired confidence. We show that several other voting rules (e.g., Borda count) require a strictly higher, polynomial number of samples, while rules such as plurality and veto require an even higher, exponential number of samples.

Taking one step further and adopting a more normative viewpoint, we seek voting rules that satisfy *accuracy in the limit*: given an *infinite* number of votes, the voting rule should be guaranteed to return the correct ranking. Again, focusing on Mallows' model, we show that two wide families of voting rules — PM-c rules and position-dominance consistent (PD-c) rules — are accurate in the limit. These families are disjoint and collectively encompass most popular voting rules. Finally, we observe that this approach still produces voting rules that are tailored to a specific noise model, and thus would not provide any guarantees if the assumed noise model differs from the one that really governs the generation of noisy votes in practice. We thus seek voting rules that provide accuracy in the limit with respect to *all* noise models within a broad class of noise models. We pinpoint the exact classes of noise models (of a specific form) for which the families of PM-c and PD-c rules are accurate in the limit, and show that these are indeed broad classes of noise models, as desired. The final result in the chapter identifies an extremely broad class of noise models, and pinpoints the *unique* voting rule, which

[12] Intuitively, if a ranking is not obtained because of cycle formation, the process is restarted.

we term the *modal ranking* rule, that is accurate in the limit with respect to every noise model in this class.

Chapter 8: Robust Voting on Social Networks

Previous work in the statistical framework of voting (including the work described in Chapter 7) relies on a crucial modeling assumption: the votes are independent (conditional on the truth). This assumption is clearly satisfied in some settings, but in many other settings — especially when the voters are people — votes are likely to be correlated through social interactions. We refer to the structure of these interactions as a *social network*, interpreted in the broadest possible sense: any form of interaction qualifies for an edge. From this broad viewpoint, the structure of the social network cannot be known, and, hence, votes are correlated in an unpredictable way. Inspired by the robustness approach from Chapter 7, our goal in this chapter is to model the generation of noisy rankings on a social network given a ground truth, and identify voting rules that are accurate in the limit with respect to any network structure and (almost) any choice of model parameters.

Our model is inspired from the *independent conversations model* due to Conitzer [63]. In our model, we assume that each alternative a has a true quality μ_a. The result of an independent conversation on an edge is a noisy quality estimate for each alternative a sampled from a Gaussian distribution with mean μ_a. Each voter assigns a weight to each incident edge, and computes an aggregate quality estimate for each alternative a by taking a weighted average of the noisy quality estimates of a on the incident edges. The voter submits a ranking of the alternatives by their aggregate quality estimates.

We analyze the performance of PM-c rules, PD-c rules, and the modal ranking rule introduced in Chapter 7. Under a mild condition on the weights placed by the voters on their incident edges, we show that all PM-c rules, an important subset of PD-c rules, and the modal ranking rule are accurate in the limit when all the Gaussian distributions have equal variance. However, when the Gaussians can have unequal variance, many PD-c rules and the modal ranking rule are no longer accurate in the limit, whereas all PM-c rules stay accurate in the limit. Therefore, PM-c rules exhibit qualitatively more robustness than PD-c rules and the modal ranking rule.

Our work is similar in flavor to the work of Conitzer [62, 63]. However, his latter work [63] only supports two alternatives, and his former work [62] investigates a model in which the (known) network structure is essentially irrelevant in that the maximum likelihood estimator does not depend on it. A bit further afield, there is a large body of work that studies the diffusion of opinions, votes, technologies, or products (but not ranked estimates) in a social network. An especially pertinent example is the work of Mossel et al. [147], where at each time step voters adopt the most popular opinion among their neighbors, and at some point opinions are aggregated via the plurality rule. Other popular diffusion models that are similar to our work include the *independent cascade model*, the *linear threshold model*, and the DeGroot model [72].

Chapter 9: A Worst-Case Approach to Voting

In the statistical framework of voting, the MLE approach is specific to a noise model, and that noise model — even if it exists for a given setting — may be difficult to pin down [140]. The robustness approach proposed in Chapters 7 and 8 alleviates this issue, but it may potentially require an infinite amount of information. Further, one limitation that applies to the entire statistical framework itself is that it can only provide compelling guarantees when the number of votes is large. Intuitively, however, if we receive a small number of votes that are very close to the ground truth, their aggregation should be accurate as well.

In this chapter, we propose the first worst-case (in other words, adversarial) approach to aggregating noisy votes that achieves these desired properties. Instead of assuming probabilistic noise, we assume a known upper bound t on the "average noise" (measured by distance from the ground truth, according to a distance metric d) in the input votes, and allow the input votes to be *adversarial* subject to the upper bound. And because it is not always possible to recover the ground truth, we wish to recover a ranking that is guaranteed to be at distance at most k from the ground truth ranking, and seek bounds on k in terms of t and the number of votes n.

We emphasize that in potential application domains there is no adversary that actively inserts errors into the votes; we choose an adversarial error model to be able to correct errors even in the *worst case*. This style of worst-case analysis is prevalent in many branches of computer science, e.g., in the analysis of online algorithms [32], and in machine learning [26, 117].

Our approach is very closely related to the vast literature on error-correcting codes that uses permutations [see, e.g., 17, and the references therein], and especially to list decoding of error-correcting codes [see, e.g., 107]. See Chapter 9 for details. Within social choice theory, our model is reminiscent of distance rationalizability with strongly unanimous consensus [143], in which one finds the ranking closest to the profile. In our approach, however, we look at, not only the closest, but all rankings up to an average distance of t from the given profile — as they are all plausible ground truths — and return a single ranking that is at distance at most k from all such rankings.

Our theoretical results are threefold. First, we observe that for any distance metric, one can always recover a ranking that is at distance at most $2t$ from the ground truth, i.e., $k \leqslant 2t$. Subject to polynomial running time, we show a weaker bound $k \leqslant 3t$. Second, we complement the upper bounds by providing a universal lower bound of (roughly) $k \geqslant t/2$ that holds for every distance metric. By imposing an extremely mild assumption on the distance metric, we can improve the lower bound to (roughly) $k \geqslant t$. In addition, we consider the four most popular distance metrics used in the social choice literature, and, for each, prove a tight lower bound of (roughly) $k \geqslant 2t$. Third, in practice the upper bound t may not be highly accurate. We provide theoretical performance guarantees in cases where the upper bound t is an underestimate or overestimate of the tightest upper bound.

Finally, we also test our worst-case-optimal voting rules against many well-known voting rules on two real-world datasets [140], and show that the worst-case optimal rules exhibit superior performance as long as the given error bound t is a reasonable overestimate of the tightest upper bound.

1.4 Prerequisites

This thesis requires an undergraduate-level mathematical knowledge of topics such as discrete mathematics, graph theory, algebra, and probability theory, and a graduate-level knowledge of the theory of computer science including topics such as basic complexity theory, approximation algorithms, randomized algorithms, and optimization.

In particular, the thesis is self-contained with respect to its economic aspects: all the required economic concepts are introduced in the relevant chapters.

1.5 Bibliographic Notes

The research presented in this thesis is based on joint work with many co-authors, as described below. In each work, I am either the primary contributor or one of two equal contributors.

In Part I (fair division), Chapter 2 is based on joint work with Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, and Junxing Wang [50]. Chapter 3 is based on joint work with David Kurokawa and Ariel D. Procaccia [126]. Chapter 4 is based on joint work with David C. Parkes and Ariel D. Procaccia [165]. Chapter 5 is based on joint work with Ian Kash and Ariel D. Procaccia [165].

In Part II (social choice theory), Chapter 6 is based on joint work with Ioannis Caragiannis, Swaparava Nath, and Ariel D. Procaccia [51]. Chapter 7 is based on joint work with Ioannis Caragiannis and Ariel D. Procaccia [49, 52]. Chapter 8 is based on joint work with Ariel D. Procaccia and Eric Sodomka [181]. Finally, Chapter 9 is based on joint work with Ariel D. Procaccia and Yair Zick [182].

1.5.1 Excluded Research

A significant portion of my work during my PhD studies has been excluded from this thesis to keep the thesis succinct and easy to read. While some of these works falls in the fair division and social choice realms, many do not. The excluded research includes:

- Work on computational social choice: uncertainty [173], the maximum likelihood approach [82, 172, 179], geometry of voting [128], and applications to multi-agent systems [115].
- Work on security games: multi-defender coordination [114].
- Work on cooperative game theory: agent failures [14, 15, 16] and the structure of synergies [180].
- Work on recommendation systems: false-name-proof recommendations on social networks [42].
- Work on machine learning: truthful univariate estimation [53].
- Work on prediction markets: theoretical analysis of long-run wealth and market price in a closed prediction market [120].

Part I

Fair Division

Chapter 2

Allocating Goods and Maximum Nash Welfare

2.1 Introduction

In this chapter, we are interested in the problem of fairly allocating *indivisible* goods, such as jewelry or artworks. But to better understand the context for our work, let us start with an easier problem: fairly allocating *divisible* goods. Specifically, let there be m homogeneous divisible goods, and n players with *linear* valuations over these goods, that is, if player i receives an x_{ig} fraction of good g, her value is $v_i(\mathbf{x}_i) = \sum_g x_{ig} v_i(g)$, where $v_i(g)$ is her non-negative value for the (entire) good g alone.

The question, of course, is what fraction of each good to allocate to each player; and it has an elegant answer, given more than four decades ago by Varian [195]. Under his *competitive equilibrium from equal incomes (CEEI)* solution, all players are endowed with an equal budget, say \$1 each. The equilibrium is an allocation coupled with (virtual) prices for the goods, such that each player buys her favorite bundle of goods for the given prices, and the market clears (all goods are sold). One formal way to argue that this solution is fair is through the compelling notion of *envy freeness* [93]: Each player weakly prefers her own bundle to the bundle of any other player. This property is obviously satisfied by CEEI, as each player can afford the bundle of any other player, but instead chose to buy her own bundle.

While the CEEI solution may seem technically unwieldy at first glance, it always exists, and, in fact, has a very simple structure in the foregoing setting: the CEEI allocations (which are what we care about, as the prices are virtual) exactly coincide with allocations $\mathbf{x}$ that maximize the *Nash social welfare* $\prod_i v_i(\mathbf{x}_i)$ [6, Volume 2, Chapter 14]. Consequently, a CEEI allocation can be computed in polynomial time via the convex program of Eisenberg and Gale [80].

Let us now revisit our original problem — that of allocating *indivisible* goods, under *additive valuations*: the utility of a player for her bundle of goods is simply the sum of her values for the individual goods she receives. This is an inhospitable world where central fairness notions like envy freeness cannot be guaranteed (just think of a single indivisible good and two players). Needless to say, the existence of a CEEI allocation is no longer assured.

Nevertheless, the idea of maximizing the Nash social welfare (that is, the product of utilities) seems natural in and of itself [59, 184]. Informally, it hits a sweet spot between Bentham's utilitarian notion of social welfare — maximize the sum of utilities — and the egalitarian notion of Rawls — maximize the minimum utility. Moreover, this solution is *scale-free*, in the sense that scaling a player's valuation function would not change the outcome [151]. But, when the maximum Nash welfare solution is wrenched from the world of divisible goods, it seems to lose its potency. Or does it?

Our goal in this chapter is to demonstrate the "unreasonable effectiveness" [202] — or unreasonable fairness, if you will — of the *maximum Nash welfare (MNW)* solution, even when the goods are indivisible. We wish to convince the reader that

> *... the MNW solution exhibits an elusive combination of fairness and efficiency properties, and can be easily computed in practice. It provides the most practicable approach to date — arguably, the ultimate solution — for the division of indivisible goods under additive valuations.*

2.1.1 Real-World Connections and Implications

Our quest for fairer algorithms is part of the growing body of work on practical applications of computational fair division [1, 44, 99, 126, 175]. We are especially excited about making a real-world impact through *Spliddit* (`www.spliddit.org`), a not-for-profit fair division website [104]. Since its launch in November 2014, the website has attracted more than 60,000 users. The motto of Spliddit is *provably fair solutions*, meaning that the solutions obtained from each of the website's five applications satisfy guaranteed fairness properties. These properties are carefully explained to users, thereby helping users understand why the solutions are fair and increasing the likelihood that they would be adopted (in contrast, trying to explain the algorithms themselves would be much trickier).

One of Spliddit's five applications is allocating goods. In our view it is the hardest problem Spliddit attempts to solve, and the previous solution left something to be desired; here is how it worked. First, to express their preferences, users simply need to divide 1000 points between the goods. This simple elicitation process relies on the additivity assumption, and is the reason why, in our view, it is indispensable in practical applications. Given these inputs, the algorithm considered three levels of fairness: envy freeness, proportionality, and *maximin share guarantee*. The algorithm found the highest feasible level of fairness, and subject to that, maximized utilitarian social welfare. Importantly, a *maximin share allocation* may not exist, but a $(2/3)$-approximation thereof is always feasible [175]. This allowed Spliddit to provide a *provable* fairness guarantee for indivisible goods. That said, a (full) maximin share allocation can always be found in practice [36, 127].

While the algorithm generally provided good solutions, it was highly discontinuous, and its direct reliance on the maximin share alone — when envy freeness and proportionality cannot be obtained — sometimes led to nonintuitive outcomes. For example, consider this excerpt from an email from sent by a Spliddit user on January 7, 2016:

> *"Hi! Great app* :) *We're 4 brothers that need to divide an inheritance of 30+ furniture items. This will save us a fist fight* ;) *I played around with the demo app and it seems there are non-optimal results for at least two cases where everyone distributes the same amount of value onto the same goods. ... Try 3 people, 5 goods, with everyone placing 200 on every good. ... [This] case gives 3 to one person and 1 to each of the others. Why is that?"*

The answer to the user's question is that envy freeness and proportionality are infeasible in the example, so the algorithm sought a maximin share allocation. In every partition of the five goods into three bundles there is a bundle with at most one good (worth 200 points), hence the maximin share guarantee of each player is 200 points. Therefore, giving three goods to one player and one good to each of the others indeed maximizes utilitarian social welfare subject to giving each player her maximin share guarantee. Note that the MNW solution produces the intuitively fair allocation in this example (two players receive two goods each, one player receives one good).

Based on the results described in this chapter, we firmly believe that the MNW solution is superior to the incumbent algorithm for allocating goods and to every other approach we know of. It has been deployed on Spliddit on May 24, 2016.

2.2 The Model

Let $[k] \triangleq \{1, \ldots, k\}$. We wish to divide a set of *indivisible goods* $\mathcal{M}$ (with $m = |\mathcal{M}|$) among a set of *players* $\mathcal{N} = [n]$. As we described in Section 1.2.1, this is without loss of generality because a divisible good can always be modeled as a collection of k indivisible goods, and the model becomes accurate as k goes to infinity. See Section 2.7 for a discussion on how our methods, its guarantees, and our implementation extend to the case where some of the goods are divisible (see Section 2.7).

Next, each player i is endowed with a *valuation* function $v_i : 2^{\mathcal{M}} \to \mathbb{R}_{\geqslant 0}$ such that $v_i(\emptyset) = 0$ and $v_i(X) \leqslant v_i(Y)$ for $X \subseteq Y \subseteq \mathcal{M}$. Recall from Section 1.2.1 that *additive* valuations satisfy $v_i(S) = \sum_{g \in S} v_i(\{g\})$ for all $S \subseteq \mathcal{M}$. To simplify notation, we write $v_i(g)$ instead of $v_i(\{g\})$ for a good $g \in \mathcal{M}$. While all deployed implementations of fair division methods for indivisible goods — including Adjusted Winner [37] and the algorithm implemented on Spliddit (see Section 2.1.1) — rely on additive valuations, we study more expressive submodular valuations in Section 2.3.1.

For $S \subseteq \mathcal{M}$ and $k \in \mathbb{N}$, let $\Pi_k(S)$ denote the set of ordered partitions of S into k bundles. A feasible allocation $\boldsymbol{A} = (A_1, \ldots, A_n) \in \Pi_n(\mathcal{M})$ is a partition of the goods that assigns a subset A_i of goods to each player i, and gives player i *utility* $v_i(A_i)$. Finally, recall from Section 1.2.1 that our goal is to find a feasible allocation that satisfies fairness desiderata such as envy-freeness (EF), envy-freeness up to one good (EF1), proportionality, and the maximin share guarantee (MMS), and efficiency desiderata such as Pareto optimality (PO).

2.3 Maximum Nash Welfare is EF1 and PO

The gold standard of fairness — envy freeness (EF) — cannot be guaranteed in the context of indivisible goods. In contrast, envy freeness up to one good (EF1) is surprisingly easy to achieve under additive valuations.

Indeed, under the *draft* mechanism, the goods are allocated in a round-robin fashion: each of the players $1, \ldots, n$ selects her most preferred good in that order, and we repeat this process until all the goods have been selected. To see why this allocation is EF1, consider some player $i \in \mathcal{N}$. We can partition the sequence of choices $1, \ldots, i-1, i, i+1, \ldots, n, 1, \ldots, i-1, \ldots$ into *phases* $i, \ldots, i-1$, each starting when player i makes a choice, and ending just before she makes the next choice. In each phase, i receives a good that she (weakly) prefers to each of the $n-1$ goods selected by subsequent players. The only potential source of envy is the goods selected by players $1, \ldots, i-1$ before the beginning of the first phase (that is, before i ever chose a good); but there is at most one such good per player $j \in [i-1]$, and removing that good from the bundle of j eliminates any envy that i might have had towards j.

However, it is clear that the allocation returned by the draft mechanism is not guaranteed to be Pareto optimal. One intuitive way to see this is that the draft outcome is highly constrained, in that all players receive almost the same number of goods; and mutually beneficial swaps of one good in return for multiple goods are possible.

Is there a different approach for generating allocations that are EF1 and PO? Surprisingly, several natural candidates fail. For example, maximizing the utilitarian welfare (the sum of utilities to the players) or the egalitarian welfare (the minimum utility to any player) is not EF1 (see Example A.2 in Appendix A.3). Interestingly, maximizing these objectives subject to the constraint that the allocation is EF1 violates PO (see Example A.3 in Appendix A.3, which was generated through computer simulations).

An especially promising idea — which was our starting point for the research reported herein — is to compute a CEEI allocation assuming the goods are divisible, and then to come up with an intelligent rounding scheme to allocate each good to one of the players who received some fraction of it. The hope was that, because the CEEI allocation is known to be EF for divisible goods [195], some rounding scheme, while inevitably violating EF, will only create envy up to one good, i.e., will still satisfy EF1. But we found a counterexample in which *every* rounding of the "divisible CEEI" allocation violates EF1; this is presented as Example A.1 in Appendix A.3.

As mentioned earlier, for divisible goods a CEEI allocation maximizes the Nash welfare. And, although a CEEI allocation may not exist for indivisible goods, one can still maximize the Nash welfare over all feasible allocations. Strikingly, this solution, which we refer to as the *maximum Nash welfare (MNW)* solution, achieves both EF1 and PO.

Definition 2.1 (The MNW solution)**.** The Nash welfare of allocation $\boldsymbol{A} \in \Pi_n(\mathcal{M})$ is defined as $\texttt{NW}(\boldsymbol{A}) = \prod_{i \in \mathcal{N}} v_i(A_i)$. Given valuations $\{v_i\}_{i \in \mathcal{N}}$, the MNW solution selects an allocation $\boldsymbol{A}^{\text{MNW}}$ maximizing the Nash welfare among all feasible allocations, i.e.,

$$\boldsymbol{A}^{\text{MNW}} \in \arg\max_{\boldsymbol{A} \in \Pi_n(\mathcal{M})} \texttt{NW}(\boldsymbol{A}).$$

If it is possible to achieve positive Nash welfare (i.e., provide positive utility to every player simultaneously), any Nash-welfare-maximizing allocation can be selected. In the special case that every feasible allocation has zero Nash welfare (i.e., it is impossible to provide positive utility to every player simultaneously), we find a largest set of players to which we can simultaneously provide positive utility, and select an allocation to these players maximizing their product of utilities. While this edge case is highly unlikely to appear in practice, it must be handled carefully to retain the solution's attractive fairness and efficiency properties. We say that an allocation is a *maximum Nash welfare (MNW)* allocation if it can be selected by the MNW solution. The MNW solution is formally specified as Algorithm 9 in Appendix A.1.

We are now ready to state our first result, which is relatively simple yet, we believe, especially compelling.

Theorem 2.1. *Every MNW allocation is envy free up to one good (EF1) and Pareto optimal (PO) for additive valuations over indivisible goods.*

Proof. Let $\boldsymbol{A}$ denote an MNW allocation. First, let us assume $\texttt{NW}(\boldsymbol{A}) > 0$. Pareto optimality of $\boldsymbol{A}$ holds trivially because an alternative allocation that increases the utility to some players without decreasing the utility to any player would increase the Nash welfare, contradicting the optimality of the Nash welfare under $\boldsymbol{A}$. Suppose, for contradiction, that $\boldsymbol{A}$ is not EF1, and that player i envies player j even after removing any single good from player j's bundle.

Let $g^* = \arg\min_{g \in A_j, v_i(g)>0} v_j(g)/v_i(g)$. Note that g^* is well-defined because player i envying player j implies that player i has a positive value for at least one good in A_j. Let $\boldsymbol{A'}$ denote the allocation obtained by moving g^* from player j to player i in $\boldsymbol{A}$. We now show that $\texttt{NW}(\boldsymbol{A'}) > \texttt{NW}(\boldsymbol{A})$, which gives the desired contradiction as the Nash welfare is optimal under $\boldsymbol{A}$. Specifically, we show that $\texttt{NW}(\boldsymbol{A'})/\texttt{NW}(\boldsymbol{A}) > 1$. The ratio is well-defined because we assumed $\texttt{NW}(\boldsymbol{A}) > 0$.

Note that $v_k(A'_k) = v_k(A_k)$ for all $k \in \mathcal{N} \setminus \{i,j\}$, $v_i(A'_i) = v_i(A_i) + v_i(g^*)$, and $v_j(A'_j) = v_j(A_j) - v_j(g^*)$. Hence,

$$\frac{\texttt{NW}(\boldsymbol{A'})}{\texttt{NW}(\boldsymbol{A})} > 1 \Leftrightarrow \left[1 - \frac{v_j(g^*)}{v_j(A_j)}\right] \cdot \left[1 + \frac{v_i(g^*)}{v_i(A_i)}\right] > 1 \Leftrightarrow \frac{v_j(g^*)}{v_i(g^*)} \cdot \left[v_i(A_i) + v_i(g^*)\right] < v_j(A_j), \quad (2.1)$$

where the last transition follows using simple algebra. Due to our choice of g^*, we have

$$\frac{v_j(g^*)}{v_i(g^*)} \leqslant \frac{\sum_{g \in A_j} v_j(g)}{\sum_{g \in A_j} v_i(g)} = \frac{v_j(A_j)}{v_i(A_j)}. \quad (2.2)$$

Because player i envies player j even after removing g^* from player j's bundle, we have

$$v_i(A_i) + v_i(g^*) < v_i(A_j). \quad (2.3)$$

Multiplying Equations (2.2) and (2.3) gives us the desired Equation (2.1).

Let us now address the special case where $\texttt{NW}(\boldsymbol{A}) = 0$. Let S denote the set of players to which the solution gives positive utility. Then, by the definition of the MNW solution (see Algorithm 9), S is a largest set of players to which one can provide positive utility. Pareto optimality of $\boldsymbol{A}$ now follows easily. An alternative allocation that does not reduce the utility to any player (and thus gives positive utility to each player in S) cannot give positive utility to any player in $\mathcal{N} \setminus S$. It also cannot increase the utility to a player in S because that would increase the product of utilities to the players in S, which $\boldsymbol{A}$ already maximizes.

From the proof of the case of $\texttt{NW}(\boldsymbol{A}) > 0$, we already know that there is no envy up to one good among players in S because $\boldsymbol{A}$ is an MNW allocation over these players, and under $\boldsymbol{A}$ the product of utilities to the players in S is positive. Further, because players in $\mathcal{N} \setminus S$ do not receive any goods, we only need to show that player $i \in \mathcal{N} \setminus S$ does not envy player $j \in S$ up to one good. Suppose for contradiction that she does. Choose $g_j \in A_j$ such that $v_j(g_j) > 0$. Such a good exists because we know $v_j(A_j) > 0$. Because player i envies player j up to one good, we have $v_i(A_j \setminus \{g_j\}) > v_i(A_i) = 0$. Hence, there exists a good $g_i \in A_j \setminus \{g_j\}$ such that $v_i(g_i) > 0$. However, in that case moving good g_i from player j to player i provides positive utility to player i while retaining positive utility to player j (because player j still has good g_j with $v_j(g_j) > 0$). This contradicts the fact that S is a largest set of players to which one can provide positive utility. Hence, the MNW allocation $\boldsymbol{A}$ is both EF1 and PO. ■

2.3.1 General Valuations

Heretofore we have focused on the case of additive valuations. As we argued earlier, this case is crucial in practice. But it is nevertheless of theoretical interest to understand whether the guarantees extend to larger classes of combinatorial valuations.

Specifically, Theorem 2.1 states that MNW guarantees EF1 and PO. We ask whether the same guarantees can be achieved for subadditive, superadditive, submodular (a special case of subadditive), and supermodular (a special case of superadditive) valuations. The definitions of these valuation classes as well as the proofs of all the results in this section are provided in Appendix A.4. Unfortunately, we obtain a negative result for three of the four valuation classes.

Theorem 2.2. *For the classes of subadditive and supermodular (and thus superadditive) valuations over indivisible goods, there exist instances that do not admit allocations that are envy free up to one good and Pareto optimal.*

We were unable to settle this question for the class of submodular valuations. And although there exist examples with submodular valuations (see, e.g., Example A.5) in which no MNW allocation satisfies EF1, we can show that every MNW allocation satisfies a relaxation of EF1 together with PO.

Definition 2.2 (MEF1: Marginal Envy Freeness Up To One Good)**.** We say that an allocation $A \in \Pi_n(\mathcal{M})$ satisfies MEF1 if

$$\forall i, j \in \mathcal{N}, \exists g \in A_j, v_i(A_i) \geqslant v_i(A_i \cup A_j \setminus \{g\}) - v_i(A_i).$$

Note that MEF1 is strictly weaker than EF1. However, for additive valuations MEF1 coincides with EF1. Hence, Theorem 2.1 follows directly from the next result (although our direct proof of Theorem 2.1 is simpler).

Theorem 2.3. *Every MNW allocation satisfies marginal envy freeness up to one good (MEF1) and Pareto optimality (PO) for submodular valuations over indivisible goods.*

2.4 Maximum Nash Welfare is Approximately MMS

In this section, we show that the fairness properties of the MNW solution extend to an alternative fairness notion — the maximin share guarantee (which is a relaxation of envy-freeness when all goods are allocated and valuations are additive), as well as a variant thereof — in theory and practice.

2.4.1 Approximate MMS, in Theory

From a technical viewpoint, our most involved result is the following theorem.

Theorem 2.4. *Every MNW allocation is π_n-maximin share (MMS) for additive valuations over indivisible goods, where*

$$\pi_n = \frac{2}{1 + \sqrt{4n - 3}}.$$

Further, the factor π_n is tight, i.e., for every $n \in \mathbb{N}$ and $\varepsilon > 0$, there exists an instance with n players having additive valuations in which no MNW allocation is $(\pi_n + \varepsilon)$-MMS.

Before we provide a proof, let us recall that the best known approximation of the MMS guarantee — to date — is $2/3 + O(1/n)$ [175], where the bound for $n = 3$ is $3/4$. But the only known way to achieve a good bound is to build the algorithm around the MMS approximation goal [2, 175]. In contrast, the MNW solution achieves its $\pi_n = \Theta(1/\sqrt{n})$ ratio "organically", as one of several attractive properties. Moreover, in almost all real-world instances, the number of players n is fairly small. For example, on Spliddit, the average number of players is very close to 3, for which our worst-case approximation guarantee is $\pi_3 = 1/2$ — qualitatively similar to $3/4$. That said, the approximation ratio achieved on real-world instances is significantly better (see Section 2.4.3).

of Theorem 2.4. We first prove that an MNW allocation is π_n-MMS (lower bound), and later prove tightness of the approximation ratio π_n (upper bound).

Proof of the lower bound: Let $\boldsymbol{A}$ be an MNW allocation. As in the proof of Theorem 2.1, we begin by assuming $\texttt{NW}(\boldsymbol{A}) > 0$, and handle the case of $\texttt{NW}(\boldsymbol{A}) = 0$ later. Fix a player $i \in \mathcal{N}$. For a player $j \in \mathcal{N} \setminus \{i\}$, let $g_j^* = \arg\max_{g \in A_j} v_i(g)$ denote the good in player j's bundle that player i values the most. We need to establish an important lemma.

Lemma 2.1. *It holds that*

$$v_i(A_j \setminus \{g_j^*\}) \leqslant \min\left\{v_i(A_i), \frac{(v_i(A_i))^2}{v_i(g_j^*)}\right\},$$

where the RHS is defined to be $v_i(A_i)$ if $v_i(g_j^*) = 0$.

Proof. First, $v_i(A_j \setminus \{g_j^*\}) \leqslant v_i(A_i)$ follows directly from the fact that $\boldsymbol{A}$ is an MNW allocation, and is therefore EF1 (Theorem 2.1). If $v_i(g_j^*) = 0$, then we are done. Assume $v_i(g_j^*) > 0$. By the definition of an MNW allocation, moving good g_j^* from player j to player i should not increase the Nash welfare. Thus,

$$v_i(A_i \cup \{g_j^*\}) \cdot v_j(A_j \setminus \{g_j^*\}) \leqslant v_i(A_i) \cdot v_j(A_j) \;\Rightarrow\; v_j(g_j^*) \geqslant v_j(A_j) - \frac{v_i(A_i) \cdot v_j(A_j)}{v_i(A_i \cup \{g_j^*\})}. \tag{2.4}$$

Note that the RHS in the above expression is positive because $v_i(g_j^*) > 0$. Hence, we also have $v_j(g_j^*) > 0$. Similarly, moving all the goods in A_j except g_j^* from player j to player i should also not increase the Nash welfare. Hence,

$$v_i(A_i \cup A_j \setminus \{g_j^*\}) \cdot v_j(g_j^*) \leqslant v_i(A_i) \cdot v_j(A_j).$$

We conclude that

$$\begin{aligned} v_i(A_j \setminus \{g_j^*\}) &\leqslant \frac{v_i(A_i) \cdot v_j(A_j)}{v_j(g_j^*)} - v_i(A_i) \leqslant \frac{v_i(A_i) \cdot v_j(A_j)}{v_j(A_j) - \frac{v_i(A_i) \cdot v_j(A_j)}{v_i(A_i \cup \{g_j^*\})}} - v_i(A_i) \\ &= v_i(A_i) \cdot \left[\frac{1}{1 - \frac{v_i(A_i)}{v_i(A_i \cup \{g_j^*\})}} - 1\right] = v_i(A_i) \cdot \left[\frac{v_i(A_i \cup \{g_j^*\})}{v_i(g_j^*)} - 1\right] = \frac{(v_i(A_i))^2}{v_i(g_j^*)}, \end{aligned}$$

where the second transition follows from Equation (2.4). ■ (Proof of Lemma 2.1)

Now, let us find an upper bound on the MMS guarantee for player i. Recall that $\texttt{MMS}_i$ is the maximum value player i can guarantee herself if she partitions the set of goods into n bundles but receives her least valued bundle. The key intuition is that indivisibility of the goods only restricts the player in terms of the partitions she can create. That is, if some of the goods become *divisible*, it can only increase the MMS guarantee of the player as she can still create all the bundles that she could with indivisible goods.

Suppose all the goods except goods in $T = \{g_j^* : j \in \mathcal{N} \setminus \{i\}, v_i(g_j^*) > \texttt{MMS}_i\}$ become divisible. It is easy to see that in the following partition, player i's value for each bundle must be at least $\texttt{MMS}_i$: put each good in T (entirely) in its own bundle, and divide the rest of the goods into $n - |T|$ bundles of equal value to player i. Because each of the latter $n - |T|$ bundles must have value at least $\texttt{MMS}_i$ for player i, we get

$$\texttt{MMS}_i \leqslant \frac{v_i(A_i) + \sum_{j \in N \setminus \{i\}} \left(v_i(g_j^*) \cdot \mathbb{I}\left[v_i(g_j^*) \leqslant \texttt{MMS}_i\right] + v_i(A_j \setminus \{g_j^*\})\right)}{n - \sum_{j \in N \setminus \{i\}} \left[v_i(g_j^*) > \texttt{MMS}_i\right]}, \tag{2.5}$$

where $\mathbb{I}(\cdot)$ denotes the indicator function.

Next, we use the upper bound on $v_i(A_j \setminus \{g_j^*\})$ from Lemma 2.1, and divide both sides of Equation (2.5) by $v_i(A_i)$. For simplicity, let us denote $x_j = v_i(g_j^*)/v_i(A_i)$, and $\beta = \texttt{MMS}_i/v_i(A_i)$. Note that β is the reciprocal of the bound on the MMS approximation that we are interested in. Then, we get

$$\beta \leqslant \frac{1 + \sum_{j \in N \setminus \{i\}} \left(x_j \cdot \mathbb{I}\left[x_j \leqslant \beta\right] + \min\left\{1, \frac{1}{x_j}\right\}\right)}{n - \sum_{j \in N \setminus \{i\}} \mathbb{I}\left[x_j > \beta\right]}.$$

Let $f(\boldsymbol{x}; \beta)$ denote the RHS of the inequality above. Then, we can write $\beta \leqslant f(\boldsymbol{x}; \beta) \leqslant \max_{\boldsymbol{x}} f(\boldsymbol{x}; \beta)$. Note that if $\beta \leqslant 1$ then player i is already receiving her full maximin share value, which gives a (stronger than) desired MMS approximation. Let us therefore assume that $\beta > 1$. To find the maximum value of $f(\boldsymbol{x}; \beta)$ over all $\boldsymbol{x}$, let us take its partial derivative with respect to x_k for $k \in \mathcal{N} \setminus \{i\}$. Note that the function is differentiable at all points except $x_k = 1$ and $x_k = \beta$.

$$\frac{\partial f}{\partial x_k} = \begin{cases} \frac{1}{n - \sum_{j \in \mathcal{N} \setminus \{i\}} \mathbb{I}[x_j > \beta]} & \text{if } 0 \leqslant x_k < 1, \\ \frac{1 - (x_k)^{-2}}{n - \sum_{j \in \mathcal{N} \setminus \{i\}} \mathbb{I}[x_j > \beta]} & \text{if } 1 < x_k < \beta, \\ \frac{-(x_k)^{-2}}{n - \sum_{j \in \mathcal{N} \setminus \{i\}} \mathbb{I}[x_j > \beta]} & \text{if } \beta < x_k. \end{cases}$$

Note that $\partial f / \partial x_k > 0$ for $x \in (0, 1)$ and $x \in (1, \beta)$, and $\partial f / \partial x_k < 0$ for $x_k > \beta$. Further note that f is continuous at $x_k = 1$. Hence, the maximum value of f is achieved either at $x_k = \beta$ or in the limit as $x_k \to \beta^+$ (i.e., when x_k converges to β from above). Suppose the maximum is achieved when t of the x_k's are equal to β, and the other $n - t - 1$ approach β from above. Then, the value of f is

$$g(t; \beta) = \frac{1 + t \cdot \left(\beta + \frac{1}{\beta}\right) + (n - t - 1) \cdot \frac{1}{\beta}}{n - (n - t - 1)}.$$

We now have that $\beta \leqslant \max_{t \in \{0, \ldots, n-1\}} g(t; \beta)$. Note that

$$\frac{\partial g}{\partial t} = \frac{\beta - 1 - (n-1) \cdot \frac{1}{\beta}}{(t+1)^2}.$$

If $\beta = \texttt{MMS}_i / v_i(A_i) \leqslant 1/\pi_n$, we already have the desired MMS approximation. Assume $\beta > 1/\pi_n$. It is easy to check that this implies $\partial g / \partial t > 0$. Thus, the maximum value of g is achieved at $t = n - 1$, which gives $\beta \leqslant (1/n) \cdot (1 + (n-1) \cdot (\beta + 1/\beta))$, which simplifies to $\beta \leqslant 1/\pi_n$, which is a contradiction as we assumed $\beta > 1/\pi_n$.

Recall that for the proof above, we assumed $\texttt{NW}(\boldsymbol{A}) > 0$. Let us now handle the special case where an MNW allocation $\boldsymbol{A}$ satisfies $\texttt{NW}(\boldsymbol{A}) = 0$. Let S denote the set of players that receive positive utility under $\boldsymbol{A}$, where $|S| < n$. Due to the definition of an

MNW allocation (see Algorithm 9), $\boldsymbol{A}$ is an MNW allocation over the players in S. Thus, from the proof of the previous case, we know that each player in S in fact achieves at least a $\pi_{|S|}$-fraction of her $|S|$-player MMS guarantee, which is at least a π_n-fraction of her n-player MMS guarantee. Players in $\mathcal{N} \setminus S$ receive zero utility. We now show that their (n-player) MMS guarantee is also 0, which yields the required result.

Suppose a player $i \in \mathcal{N} \setminus S$ has a positive value for at least n goods in $\mathcal{M}$. Now, because these goods are allocated to at most $n-1$ players in S, at least one player $j \in S$ must have received at least two goods g_1 and g_2, both of which player i values positively. Because player j receives positive utility under $\boldsymbol{A}$ (i.e., $v_j(A_j) > 0$), it is easy to check that there exists a good $g \in \{g_1, g_2\}$ such that $v_j(A_j \setminus \{g\}) > 0$. Thus, moving good g to player i provides positive utility to player i while retaining positive utility to player j, which violates the fact that S is a largest set of players to which one can simultaneously provide positive utility. This shows that player i has positive utility for at most $n-1$ goods in $\mathcal{M}$, which immediately implies $\texttt{MMS}_i = 0$, as required.

Proof of the upper bound (tightness): We now show that for every $n \in \mathbb{N}$ and $\varepsilon > 0$, there exists an instance with n players in which no MNW allocation is $(\pi_n + \varepsilon)$-MMS. For $n = 1$, this is trivial because $\pi_1 = 1$. Hence, assume $n \geqslant 2$.

Let the set of players be $\mathcal{N} = \{1, \ldots, n\}$, and the set of goods be $\mathcal{M} = \{x\} \cup \bigcup_{i \in \{2,\ldots,n\}} \{h_i, l_i\}$. Thus, we have $m = 2n - 1$ goods. We refer to h_i's as the "heavy" goods and l_i's as the "light" goods. Let the valuations of the players for the goods be as follows. Choose a sufficiently small $\varepsilon' > 0$ (an upper bound on ε' will be determined later in the proof).

Player 1: $\quad v_1(x) = 1$, and $\forall j \in \{2, \ldots, n\}, v_1(h_j) = \dfrac{1}{\pi_n} - \varepsilon'$ and $v_1(l_j) = \pi_n - \varepsilon'$.

Player i, for i $\geqslant$ 2: $\quad v_i(h_i) = \dfrac{1}{\pi_n + 1}$, $v_i(l_i) = \dfrac{\pi_n}{\pi_n + 1}$, and $\forall g \in \mathcal{M} \setminus \{h_i, l_i\}, v_i(g) = 0$.

In particular, note that player 1 has a positive value for every good (for $\varepsilon' < \pi_n$), while for $i \geqslant 2$, player i has a positive value for only two goods: h_i and l_i. Consider the allocation $\boldsymbol{A}^*$ that assigns good x to player 1, and for every $i \in \mathcal{N} \setminus \{1\}$, allocates goods h_i and l_i to player i. We claim that $\boldsymbol{A}^*$ is the unique MNW allocation but is not $(\pi_n + \varepsilon)$-MMS.

First, note that an MNW allocation is Pareto optimal, and therefore it must allocate good x to player 1 because no other player has a positive value for x. Further, $\texttt{NW}(\boldsymbol{A}^*) > 0$, which implies that every MNW allocation must also have a positive Nash welfare. This in turn implies that an MNW allocation must assign to each player in $\mathcal{N} \setminus \{1\}$ at least one of h_i and l_i. Subject to these constraints, consider a candidate allocation $\boldsymbol{A}$.

Let p (resp. q) denote the number of players $i \in \mathcal{N} \setminus \{1\}$ that only receive good h_i (resp. l_i), and have utility $1/(\pi_n + 1)$ (resp. $\pi_n/(\pi_n + 1)$). Hence, exactly $n - 1 - p - q$ players $i \in \mathcal{N} \setminus \{1\}$ receive both h_i and l_i, and have utility 1. Player 1 receives good x, q heavy goods, and p light goods, and has utility $1 + q \cdot (1/\pi_n - \varepsilon') + p \cdot (\pi_n - \varepsilon')$. Thus, the Nash welfare of $\boldsymbol{A}$ is given by

$$\left(1+q\cdot\left(\frac{1}{\pi_n}-\varepsilon'\right)+p\cdot(\pi_n-\varepsilon')\right)\left(\frac{1}{\pi_n+1}\right)^p\left(\frac{\pi_n}{\pi_n+1}\right)^q=\frac{1+q\cdot\left(\frac{1}{\pi_n}-\varepsilon'\right)+p\cdot(\pi_n-\varepsilon')}{(1+\pi_n)^p\cdot\left(1+\frac{1}{\pi_n}\right)^q}.$$

Using binomial expansion, it is easy to show that the denominator in the final expression above is at least $1+p\cdot\pi_n+q/\pi_n$, which is never less than the numerator, and is equal to the numerator if and only if $p=q=0$. Note that $p=q=0$ indeed gives our desired allocation $\boldsymbol{A}^*$. Hence, the maximum Nash welfare of 1 is uniquely achieved by the allocation $\boldsymbol{A}^*$.

Next, let us analyze the MMS guarantee for player 1. In particular, consider the partition of the set of goods into n bundles $B_1,\ldots,B_n$ such that $B_1=\{x,l_2,\ldots,l_n\}$ and $B_i=\{h_i\}$ for all $i\in\{2,\ldots,n\}$. Note that for all $i\in\{2,\ldots,n\}$, $v_1(B_i)=1/\pi_n-\varepsilon'$. Also,

$$v_1(B_1)=1+(n-1)\cdot(\pi_n-\varepsilon')=1+(n-1)\cdot\pi_n-(n-1)\cdot\varepsilon'=\frac{1}{\pi_n}-(n-1)\cdot\varepsilon',$$

where the final equality holds because π_n is chosen precisely to satisfy the equation $1+(n-1)\cdot\pi_n=1/\pi_n$. As the MMS guarantee of player 1 is at least her minimum value for any bundle in $\{B_1,\ldots,B_n\}$, we have $\mathtt{MMS}_1\geqslant 1/\pi_n-(n-1)\cdot\varepsilon'$. In contrast, under the MNW allocation $\boldsymbol{A}^*$ we have $v_1(A_1)=1$. Thus, the MMS approximation ratio on this instance is at most $1/(1/\pi_n-(n-1)\cdot\varepsilon')$. It is easy to check that for driving this ratio below $\pi_n+\varepsilon$, it is sufficient to set

$$\varepsilon'<\min\left\{\pi_n,\frac{\varepsilon}{(n-1)\cdot\pi_n\cdot(\pi_n+\varepsilon)}\right\}.$$

This completes the entire proof. ■ (Proof of Theorem 2.4)

A striking aspect of the proof of Theorem 2.4 is that, at first glance, the lower bound of π_n seems very loose. For example, key steps in the proof involve the derivation of an upper bound on the MMS guarantee of player i by assuming that some of the goods are divisible, and the maximization of the function $f(\cdot)$ over an unrestricted domain. Yet the ratio π_n turns out to be completely tight.

2.4.2 Approximate Pairwise MMS, in Theory

Adding to the conceptual arguments in favor of Theorem 2.4 (see the discussion just after the theorem statement), we note that it also has interesting implications. Let us first define a novel fairness property:

Definition 2.3 (α-Pairwise Maximin Share Guarantee)**.** We say that an allocation $\boldsymbol{A}\in\Pi_n(\mathcal{M})$ is an α-pairwise maximin share (MMS) allocation if

$$\forall i,j\in\mathcal{N}, v_i(A_i)\geqslant\alpha\cdot\max_{\boldsymbol{B}\in\Pi_2(A_i\cup A_j)}\min\{v_i(B_1),v_i(B_2)\}.$$

We simply say that $\boldsymbol{A}$ is pairwise MMS if it is 1-pairwise MMS. Note that the pairwise MMS guarantee is similar to the MMS guarantee, but instead of player i partitioning the set of all items into n bundles, she partitions the combined bundle of herself and another player into two bundles, and receives the one she values less. Although neither the pairwise MMS guarantee nor the MMS guarantee imply the other, it can be shown that a pairwise MMS allocation is $(1/2)$-MMS (see Theorem A.2 in Appendix A.5).

We do not know whether a pairwise MMS allocation always exists (under the constraint that all goods must be allocated). In fact, there is an even more tantalizing and elusive fairness notion that is strictly weaker than pairwise MMS, but strictly stronger than EF1 (see Theorem A.2 in Appendix A.5, which, in particular, implies that pairwise MMS is stronger than EF1).

Definition 2.4 (EFX: Envy freeness up to the Least Valued Good)**.** We say that an allocation $\boldsymbol{A} \in \Pi_n(\mathcal{M})$ is envy free up to the least (positively) valued good if

$$\forall i, j \in \mathcal{N}, \forall g \in A_j : v_i(g) > 0, v_i(A_i) \geqslant v_i(A_j \setminus \{g\}).$$

While EF1 requires that player i not envy player j after the removal of player i's *most valued* good from player j's bundle, EFX requires that this no-envy condition would hold even after the removal of player i's *least positively valued* good from player j's bundle. Despite significant effort, we were not able to settle the question of whether an EFX allocation always exists (assuming all goods must be allocated), and leave it as an enigmatic open question.

Given this motivation for the pairwise MMS notion, it is interesting that our next result directly translates the MMS approximation bound of Theorem 2.4 into a pairwise MMS approximation. The proof of the result is in Appendix A.5.

Corollary 2.1. *Every MNW allocation is Φ-pairwise MMS, where Φ is the golden ratio conjugate, i.e., $\Phi = (\sqrt{5}-1)/2 \approx 0.618$. Further, the factor Φ is tight, i.e., for every $n \in \mathbb{N}$ and $\varepsilon > 0$, there exists an instance with n players having additive valuations in which no MNW allocation is $(\Phi + \varepsilon)$-pairwise MMS.*

2.4.3 Approximate MMS and Pairwise MMS, in Practice

Theorem 2.4 and Corollary 2.1 show that the MNW solution is *guaranteed* to be π_n-MMS and Φ-pairwise MMS. We now evaluate it on these benchmarks (which, we reiterate, it is *not* designed to optimize) using real-world data. Specifically, we use 1281 instances created so far through Spliddit's "divide goods" application. The number of players in these instances ranges from 2 to 10, and the number of goods ranges from 3 to 93. Figures 2.1(a) and 2.1(b) show the histograms of the MMS and pairwise MMS approximation ratios, respectively, achieved by the MNW solution on these instances.

Most importantly, observe that the MNW solution provides every player her full MMS (resp. pairwise MMS) guarantee, i.e., achieves the ideal 1-approximation, in more than 95% (resp. 90%) of the instances. Further, in contrast to the tight worst-case ratios of $\pi_n = \Theta(1/\sqrt{n})$ and $\Phi \approx 0.618$, the MNW solution achieves a ratio of at least 3/4 for both properties on *all* the real-world instances in our dataset.

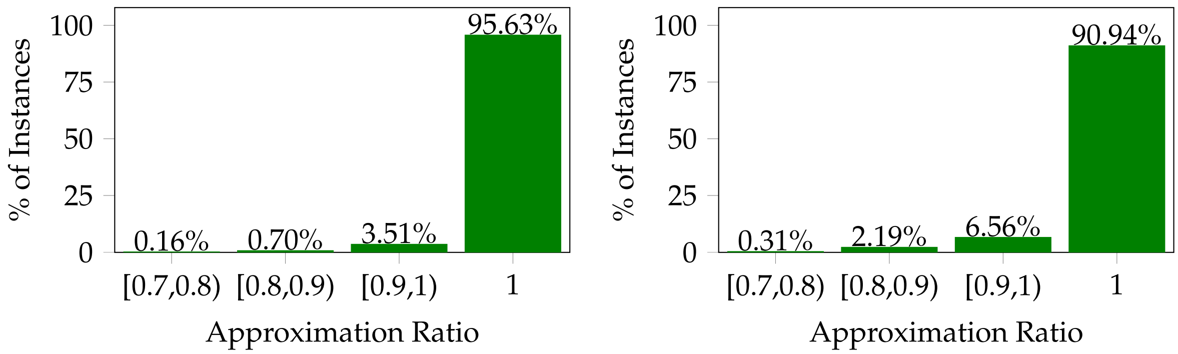

(a) MMS Approximation

(b) Pairwise MMS Approximation

Figure 2.1: MMS and Pairwise MMS approximation of the MNW solution on real-world data from Spliddit.

2.5 Implementation

It is known that computing an exact MNW allocation is $\mathcal{NP}$-hard even for 2 players with identical additive valuations, due to a simple reduction from the $\mathcal{NP}$-hard problem PARTITION [155, 184] (in fact, as we later describe, the problem is strongly $\mathcal{NP}$-hard). Our goal in this section is to develop a fast implementation of the MNW solution, despite this obstacle. An existing approach to maximizing the Nash welfare [157] iteratively modifies an initial allocation to improve the Nash welfare at each step, but may return a local maximum that does not provide any fairness or efficiency guarantees. Instead, we use integer programming to find the global optimum in a scalable way. Note that most real-world instances are relatively small, but response time can be crucial. For example, Spliddit has a demo mode, where users expect almost instantaneous results. Moreover, some instances are actually very large, as we discuss below.

Let us begin by recalling that the first step in computing an MNW allocation is to find a largest set of players S that can be given positive utility simultaneously. In Appendix A.1, we show that S can be computed easily by finding a maximum cardinality matching in an appropriate bipartite graph. The problem then reduces to computing an MNW allocation to the players in S. Hereinafter, we focus on this reduced problem. Thus, without loss of generality we can assume that for the given set of players $\mathcal{N}$, an MNW allocation will achieve positive Nash welfare.

Figure 2.2 shows a simple mathematical program for computing an MNW allocation. The binary variable $x_{i,g}$ denotes whether player i receives good g. Subject to feasibility constraints, the program maximizes the sum of log of players' utilities, or, equivalently, the Nash welfare. Note that this is a discrete optimization program with a nonlinear objective, which is typically very hard to solve.

Fortunately, we can leverage some additional properties of the problem that arise in practice. Specifically, on Spliddit, users are required to submit integral additive valuations by dividing 1000 points among the goods. This in turn ensures that the utilities to the players will also be integral, and not more than 1000. In theory, this does

$$\begin{aligned} \text{Maximize } & \textstyle\sum_{i\in\mathcal{N}} \log\left(\sum_{g\in\mathcal{M}} x_{i,g}\cdot v_i(g)\right) \\ \text{subject to } & \textstyle\sum_{i\in\mathcal{N}} x_{i,g} = 1, \forall g \in \mathcal{M} \\ & x_{i,g} \in \{0,1\}, \forall i \in \mathcal{N}, g \in \mathcal{M}. \end{aligned}$$

Figure 2.2: Nonlinear discrete optimization program

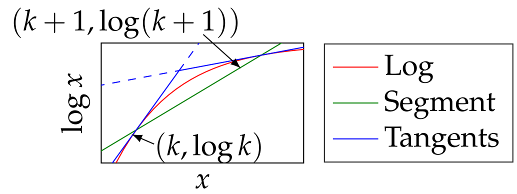

Figure 2.3: The log function and its approximations

$$\begin{aligned} \text{Maximize } & \textstyle\sum_{i\in\mathcal{N}} W_i \\ \text{subject to } & W_i \leqslant \log k + \left[\log(k+1) - \log k\right] \\ & \qquad \times \left[\textstyle\sum_{g\in\mathcal{M}} x_{i,g}\cdot v_i(g) - k\right], \\ & \qquad \forall i \in \mathcal{N}, k \in \{1,3,\ldots,999\} \\ & \textstyle\sum_{g\in\mathcal{M}} x_{i,g}\cdot v_i(g) \geqslant 1, \quad \forall i \in \mathcal{N} \\ & \textstyle\sum_{i\in\mathcal{N}} x_{i,g} = 1, \quad \forall g \in \mathcal{M} \\ & x_{i,g} \in \{0,1\}, \qquad \forall i \in \mathcal{N}, g \in \mathcal{M}. \end{aligned}$$

Figure 2.4: MILP using segments on the log curve

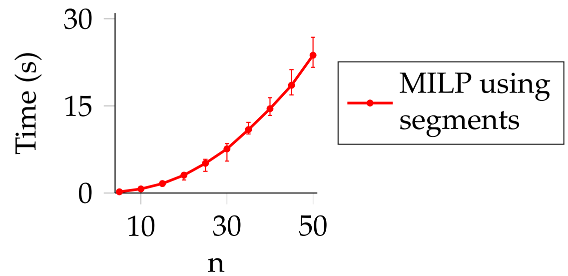

Figure 2.5: Running time of our MNW implementation

not help us: due to a known reduction from a strongly $\mathcal{NP}$-complete problem — Exact Cover by 3-Sets (X3C) — to the problem of computing an MNW allocation [155], we cannot hope for a pseudopolynomial-time algorithm (i.e., a polynomial-time algorithm for Spliddit-like valuations). In practice, however, this structure of the valuations can be leveraged to convert the non-linear objective into a linear objective: $\sum_{i\in\mathcal{N}} \sum_{t=2}^{1000} (\log t - \log(t-1)) \cdot U_{i,t}$, where $U_{i,t} = \mathbb{I}[\sum_{g\in\mathcal{M}} x_{i,g} \cdot v_i(g) \geqslant t]$ for player $i \in \mathcal{N}$ and $t \in [1000]$ is an additional variable that can be encoded using two linear constraints. However, the introduction of $1000 \cdot n$ additional binary variables makes this approach impractical even for fairly small instances.

We therefore propose an alternative approach that introduces merely n continuous variables and, crucially, no integral variables. The trick is to use a continuous variable W_i denoting the log of the utility to player i, and bound it from above using a set of linear constraints such that the tightest bound at every integral point k is exactly $\log k$. This essentially replaces the log by a piecewise linear approximation thereof that has zero error at integral points. Figure 2.3 shows two such approximations of the log function (the red line): one that uses the tangent to the log curve at the point $(k, \log k)$ for each $k \in [1000]$ (the blue lines), and one that uses segments connecting points $(k, \log k)$ and $(k+1, \log(k+1))$ for each $k \in \{1, 3, \ldots, 999\}$ (the green line). Each tangent and each segment is guaranteed to be an upper bound on the log function at every integral point due to the concavity of log.[1] Importantly, note that the tightest upper bound at each *positive* integral point k is $\log k$. These transformations do not work at $k = 0$, i.e., they do not ensure $W_i = -\infty$ if player i gets zero utility. However, recall that in our subproblem each player can achieve a positive utility. Hence, we eliminate this concern by adding the constraints that each player must receive value at least 1. We employ the transformation that uses segments as it requires half as many constraints (and, incidentally, runs nearly twice as fast). Figure 2.4 shows the final mixed-integer linear program (MILP) with only n continuous and $n \cdot m$ binary variables, which is key to the practicability of this approach.

To assess how scalable our implementation is, we measure its running time on uniformly random Spliddit-like valuations, that is, uniformly random integral valuations that sum to 1000. We vary the number of players n from 5 to 50 in increments of 5, and keep the number of goods at $m = 3 \cdot n$ to match data from Spliddit, in which $m/n \approx 3$ on average. The experiments were performed on a 2.9 GHz quad-core computer with 32 GB RAM, using CPLEX to solve the MILPs. The indicator-variables-based approach failed to run within our time limit (60 seconds) even for 5 players. Figure 2.5 shows the running time (averaged over 100 simulations, with the 5th and 95th percentiles) of the MILP formulation from Figure 2.4. Satisfyingly, we can solve instances with 50 players in less than 30 seconds (whereas even the largest of the 1281 instances on Spliddit has 10 players). In fact, the algorithm solves every Spliddit instance in less than 3 seconds.

The largest real-world instance we have seen was actually reported offline by a Spliddit user. He needed to split an inheritance of roughly 1400 goods with his 9 siblings. Our implementation solves an instance of this size in roughly 15 seconds.

2.5.1 Precision Requirements

As our optimization program involves real-valued quantities (e.g., the logarithms), we must carefully set the precision level such that the optimal allocation computed up to the precision is *guaranteed* to be an MNW allocation. This is because an allocation that only approximately maximizes the Nash welfare may fail to satisfy the theoretical guarantees of an MNW allocation (Theorems 2.1 and 2.4, and Corollary 2.1).

Recall that our objective function is the log of the Nash welfare. Hence, the difference between the objective values of an (optimal) MNW allocation and any suboptimal allocation is at least $\log(1000^n) - \log(1000^n - 1) \geqslant 1/1000^n - (1/2)/1000^{2n}$, which can be captured using $O(n)$ bits of precision. This simple observation can be easily formalized to show that there exists $p \in O(n)$ such that if all the coefficients in the optimization program are computed up to p bits, and if the program is solved with p bits of precision (i.e., with an absolute error of at most 2^{-p} in the objective function), then the solution returned will indeed correspond to an MNW allocation. Crucially, p is independent of the number of goods. We expect the number of players n to be fairly small in everyday

[1] In fact, this transformation is useful in maximizing any concave function, or minimizing any convex function, and thus may be of independent interest.

fair division problems. For example, as previously mentioned, on Spliddit more than 95% of the instances for allocating indivisible goods have $n \leqslant 3$.

Nonetheless, if one's goal is solely to find an allocation that is EF1 and PO, a constant number of bits of precision would suffice. This is because capturing differences in objective values that are at least $\log(1000^2) - \log(1000^2 - 1)$ — a constant — ensures that the resulting allocation is EF1 and PO, as we show below.

1. *EF1:* Suppose the allocation is not EF1, and player i envies player j even after the removal of any single good from player j's bundle. Then, our proof of Theorem 2.1 shows that we can increase the Nash welfare by moving a specific good from player j to player i. Because this operation does not alter the utilities to all but two players, it must increase the logarithm of the Nash welfare by at least $\log(1000^2) - \log(1000^2 - 1)$, which is a contradiction because our sensitivity level is sufficient to find this improvement.
2. *PO:* Suppose the allocation is not PO. Then there exists an alternative allocation that increases the utility to at least one player without decreasing the utility to any player. This must increase the logarithm of the Nash welfare by at least $\log(1000) - \log(1000 - 1) \geqslant \log(1000^2) - \log(1000^2 - 1)$, which is again a contradiction because our sensitivity level is sufficient to find this improvement.

2.6 Related Work

The concept of envy freeness up to one good originates in the work of Lipton et al. [134]. They deal with general combinatorial valuations, and give a polynomial-time algorithm that guarantees that the maximum envy is bounded by the maximum marginal value of any player for any good; this guarantee reduces to EF1 in the case of additive valuations. However, in the additive case, EF1 alone can be achieved by simply allocating the goods to players in a round-robin fashion, as we discuss below. The algorithm of Lipton et al. [134] does not guarantee additional properties.

Budish [44] introduces the concept of *approximate CEEI*, which is an adaptation of CEEI to the setting of indivisible goods (among other contributions in this beautiful paper, he also introduces the notion of maximin share guarantee). He shows that an approximate CEEI exists and (approximately) guarantees certain properties. The approximation error goes to zero when the number of goods is fixed, whereas the number of players, as well as the number of copies of each good, go to infinity. His approach is practicable in the MBA course allocation setting, which motivates his work — there are many students, many seats in each course, and relatively few courses. But it does not give useful guarantees for the type of instances we encounter on Spliddit, where the number of players is small, and there is typically one copy of each good.

From an algorithmic perspective, Ramezani and Endriss [184] show that maximizing Nash welfare is $\mathcal{NP}$-hard under certain combinatorial bidding languages (including, under additive valuations). Cole and Gkatzelis [59] give a constant-factor, polynomial-time approximation under additive valuations (to be precise, their objective function is the geometric mean of the utilities).[2] Lee [130] shows that the problem is APX-hard, that is, a constant-factor approximation is the best one can hope for.

When there are only two players, compelling approaches for allocating goods are available. In fact, Spliddit currently handles this case separately, via the Adjusted Winner algorithm [37]. The shortcoming of Adjusted Winner is that it usually has to split one of the goods between the two players. Adjusted Winner can be interpreted as a special case of the Egalitarian Equivalent rule of Pazner and Schmeidler [166], which is defined for any number of players. For $n > 2$ players, it may need to split all the goods, that is, it is impractical to apply it to indivisible goods.

Let us briefly mention two additional models for the division of indivisible goods. First, some papers assume that the players express *ordinal* preferences (i.e., a ranking) over the goods [13, 38]. This assumption (arguably) does not lead to crisp fairness guarantees — the goal is typically to design algorithms that compute fair allocations if they exist. Second, it is possible to allow randomized allocations [29, 30, 45]; this is hardly appropriate for the cases we find on Spliddit in which the outcome is used only once.

Finally, it is worth noting that the idea of maximizing the product of utilities was studied by Nash [154], in the context of his classic bargaining problem. This is why this notion of social welfare is named after him. In the networking community, the same solution goes by the name of *proportional fairness*, due to another property that it satisfies when goods are divisible [118]: when switching to any other allocation, the total percentage gains for players whose utilities increased sum to at most the total percentage losses for players whose utilities decreased; thus, in some sense, no such switch would be socially preferable.

2.7 Discussion

The goal of this chapter is to advocate the Maximum Nash Welfare (MNW) solution for the fair allocation of goods. While it is justified by elegant fairness (EF1) and efficiency (PO) properties, these properties are not "sufficient" *in and of themselves* — they may allow undesirable outcomes (see Example A.4 in Appendix A.3). What makes the MNW solution compelling is that it provides intuitively fair outcomes, yet *organically* satisfies these formal fairness properties. Moreover, the MNW solution provides a $\Theta(1/\sqrt{n})$-approximation to the MMS guarantee (Theorem 2.4), whereas an arbitrary EF1 and PO allocation only provides a $1/n$-approximation (Theorem A.1 in Appendix A.3).

Throughout the chapter we assumed that the goods are indivisible, but our results directly extend to the case where we have a mix of divisible and indivisible goods. The MNW solution in this case can be seen as the limit of the MNW solution on the instance where each divisible good is partitioned into k indivisible goods, as k goes to infinity. Theorem 2.1 therefore implies that the MNW solution is envy free up to one *indivisible* good, that is, player i would not envy player j (who may have both divisible

[2]However, a constant-factor approximation need not satisfy any of the theoretical guarantees we establish in this chapter for the MNW solution.

and indivisible goods) if one indivisible good is removed from the bundle of j. This provides an alternative proof for envy-freeness of the MNW/CEEI solution when all goods are divisible. The results of Section 2.4 also directly go through — in fact, the proof of the MMS approximation result (Theorem 2.4) already "liquidates" some of the goods as a technical tool. Appendix A.2 outlines the modified and scalable version of the implementation described in Section 2.5, which we have deployed on Spliddit, that can allocate a mix of divisible and indivisible goods.

It is remarkable that when all goods are divisible, three seemingly distinct solution concepts — the MNW solution, the CEEI solution, and proportional fairness (PF) — coincide. This is certainly not the case for indivisible goods: while a CEEI solution and a PF solution may not exist, the MNW solution always does. Nonetheless, our investigation revealed that even for indivisible goods, the PF solution and the MNW solution are closely related via a spectrum of solutions, which offers two advantages. First, it allows us to view the MNW solution as the optimal solution among those that lie on this spectrum and are guaranteed to exist. Second, it also gives a way to break ties — possibly even choose a unique allocation — among all MNW allocations. See Appendix A.6 for a detailed analysis. This connection between MNW and PF raises an interesting question: Is it possible to relate the MNW solution to the CEEI solution when the goods are indivisible?

Finally, we have not addressed game-theoretic questions regarding the manipulability of the MNW solution. The reason is twofold. First, there are strong impossibility results that rule out reasonable strategyproof solutions. For example, Schummer [190] shows that the only strategyproof and Pareto optimal solutions are dictatorial — which means they are maximally unfair, if you will — even when there are only two players with linear utilities over divisible goods; clearly a similar result holds for indivisible goods (at least in an approximate sense).[3] Second, we do not view manipulation as a major issue on Spliddit, because users are not fully aware of each other's preferences (they submit their evaluations in private), and — presumably, in most cases — have a very partial understanding of how the algorithm works.

[3]In theory, one can hope to circumvent this result by making manipulation computationally hard [19]. This is almost surely true (in the worst-case sense of hardness) for the MNW solution, where even computing the outcome is hard.

Chapter 3

Allocating to Strategic Agents and The Leximin Mechanism

3.1 Introduction

Our not-for-profit website Spliddit.org — which offers provably fair solutions for a range of everyday problems — has already been used by tens of thousands of people [104]. As a beautiful consequence, users are reaching out to us with questions and suggestions that have been incredibly helpful. For example, some of the users pointed out flaws in the previously deployed solution for allocating goods. As we describe in Chapter 2, this helped us design a significantly improved solution, Maximum Nash Welfare, for the case where players have *additive* valuations. More interestingly, some users are also asking for our help with specialized real-world domains that do not admit additive valuations. Designing compelling solutions for these domains has led us to investigate many interesting and difficult new questions.

This chapter presents a solution to one such question, posed by a representative of one of the largest school districts in California. Since the details are confidential, we will refer to the school district as the Pentos Unified School District (PUSD), and to the representative as Illyrio Mopatis. Mr. Mopatis contacted us in May 2014 after learning about Spliddit (and fair division, more generally) from an article in the New York Times.[1] He is tasked with the allocation of unused space (most importantly, classrooms) in PUSD's public schools to the district's charter schools, according to California's Proposition 39, which states that "public school facilities should be shared fairly among all public school pupils, including those in charter schools".[2] While the law does not elaborate on what "fairly" means, Mr. Mopatis was motivated by the belief that a provably fair solution would certainly fit the bill. He asked us to design an automated allocation method that would be evaluated by PUSD, and potentially replace the existing manual system.

[1] `http://goo.gl/Xp3omV`
[2] `http://goo.gl/bGH6dT`

To be a bit more specific, the setting consists of charter schools and facilities (public schools). Each facility has a given number of unused classrooms — its *capacity*, and each charter school has a number of required classrooms — its *demand*. In principle the classrooms required by a charter school could be split across multiple facilities, but such offers have always been declined in the past, so we assume that an agent's demand must be satisfied in a single facility (if it is satisfied at all). Other details are less important and can be abstracted away. For example, classroom size turns out to be a nonissue, and the division of time in shared space (such as the school gym or cafeteria) can be handled *ad hoc*.

Of course, to talk of fairness we must also take into account the preferences of charter schools, but preference representation is a modeling choice, intimately related to the design and guarantees of the allocation mechanism. Moreover, fairness is not our only concern: to be used in practice, the mechanism must be relatively intuitive (so it can be explained to decision makers) and computationally feasible. The challenge we address is therefore to

> *... design and implement a classroom allocation mechanism that is provably fair as well as practicable.*

3.2 Our Approach

We model the preferences of charter schools as being *dichotomous*: charter schools think of each facility as either acceptable or unacceptable. This choice is motivated by current practice: Under the 2015/2016 request form issued by PUSD, charter schools are essentially asked to indicate acceptable facilities (specifically, they are asked to "provide a description of the district school site and/or general geographic area in which the charter school wishes to locate" using free-form text). In other words, formally eliciting dichotomous preferences — by having charter schools select acceptable facilities from the list of all facilities — is similar to the *status quo*, a fact that increases the practicability of the approach.

Formally, let $N = \{1, \ldots, n\}$ denote the set of charter schools (hereinafter, *agents*), and let $M = \{1, \ldots, m\}$ denote the set of public schools (hereinafter, *facilities*). We want to design a mechanism for assigning the agents to the facilities. Each facility f has a *capacity* c_f, which is the number of *units* available at the facility (in our motivating example, each unit is a classroom). The preferences of agent i are given by a pair (d_i, F_i), where $d_i \in \mathbb{N}$ denotes the number of units demanded by agent i — or, simply, the *demand* of agent i — and $F_i \subseteq M$ denotes the set of facilities acceptable to agent i. Crucially, we assume that agent i's preferences are dichotomous in nature: the agent has utility 1 if it receives d_i units from *any single facility* $f \in F_i$ (in this case, we say agent i is assigned to facility f), and 0 otherwise. Without loss of generality, we assume that every agent i has an acceptable facility $f \in F_i$ that has sufficient capacity to meet its demand (i.e., $c_f \geqslant d_i$).[3]

[3]Agents violating this requirement cannot achieve positive utility, and can effectively be disregarded.

A deterministic allocation is a mapping $A : N \to M \cup \{0\}$, where $A_i = A(i)$ denotes the facility to which agent i is assigned (and $A_i = 0$ means agent i is not assigned to any facility). A is *feasible* if it respects the capacity constraint at each facility:

$$\forall f \in M, \sum_{i \in N : A_i = f} d_i \leqslant c_f.$$

Let $\mathcal{A}$ denote the space of all feasible deterministic allocations. Formally, the utility to agent i under a feasible deterministic allocation $A \in \mathcal{A}$ is given by

$$u_i(A_i) = \begin{cases} 1 & \text{if } A_i \in F_i \\ 0 & \text{otherwise.} \end{cases}$$

A feasible randomized allocation is simply a distribution over feasible deterministic allocations, and the utility to an agent is its expected utility under the randomized allocation. Let $\Delta(\mathcal{A})$ be the space of all feasible randomized allocations. Crucially, note that $\Delta(\mathcal{A})$ is a convex set, i.e., given randomized allocations $A, A' \in \Delta(\mathcal{A})$ and $0 \leqslant \lambda \leqslant 1$, we can construct another randomized allocation $A'' = \lambda \cdot A + (1 - \lambda) \cdot A' \in \Delta(\mathcal{A})$ that executes A with probability λ and A' with probability $1 - \lambda$. Hereinafter, an allocation is possibly randomized, unless explicitly specified otherwise.

As mentioned in Section 3.1, our setting deals with fair allocation of indivisible goods, and generalizes the classic setting of random assignment under dichotomous preferences studied by Bogomolnaia and Moulin [30]. In particular, their setting can be recovered by setting all the demands and capacities to 1 (i.e., $d_i = 1$ and $c_f = 1$ for all $i \in N$, $f \in M$), with an equal number of agents and facilities ($m = n$).

Desiderata. The fair division literature offers a slew of desirable properties. We are especially interested in four classic desiderata that have proved to be widely applicable (with applications ranging from cake cutting [169] to the division of computational resources in clusters [99, 165]), often satisfiable, and yet effective in leading to compelling mechanisms: proportionality, envy-freeness, Pareto optimality, and (group) strategyproofness. See Section 1.2.1 for their formal definitions. We use these desiderata to guide the search for a good mechanism in our setting.

Let us first consider an example illustrating the desiderata of our interest.

Example 3.1. First, let us consider a simple randomized mechanism that allocates all available units at all facilities to each agent with probability $1/n$. Clearly, the mechanism satisfies proportionality because it gives each agent utility $1/n$. The mechanism is also envy-free because each agent has an identical allocation, and thus no reason to envy any other agent. Since the mechanism operates independently of the reported preferences of the agents, the mechanism is obviously (group) strategyproof. However, the mechanism is not Pareto optimal. The reason is that the mechanism allocates all available units to an agent (with probability $1/n$) even if the agent does not require all the units. In this case, it may be possible to simultaneously satisfy another agent, thus obtaining a Pareto improvement.

Next, consider a different mechanism that always returns a deterministic allocation maximizing the number of units allocated. While this mechanism is very intuitive, we can show that it violates all the desiderata except Pareto optimality. Suppose there is a single facility with 4 available units, and two agents — namely, agents 1 and 2 — that demand 3 and 2 units, respectively. Maximizing the number of units allocated would require allocating 3 units to agent 1 and no units to agent 2. This already violates both proportionality and envy-freeness with respect to agent 2. Further, agent 2 would have a strict incentive to report a false demand of 4 units, which would lead to agent 2 receiving all 4 units from the facility. Thus, strategyproofness is also violated.

3.3 The Leximin Mechanism

A natural starting point is the seminal paper of Bogomolnaia and Moulin [30], who study the special case of our setting with unit demands and capacities, under dichotomous preferences. They propose the *leximin mechanism*, which returns a random allocation with the following intuitive property: it maximizes the lowest probability of any charter school having its demand satisfied in an acceptable facility; subject to this constraint, it maximizes the second lowest probability; and so on.

Formally, let $(u^1, u^2, \ldots, u^n)$ denote the vector of utilities sorted in non-descending order. The leximin mechanism returns the allocation that maximizes this vector in the lexicographic order; we say that this allocation is leximin-optimal. The mechanism is presented as Algorithm 1.

ALGORITHM 1: The Leximin Mechanism

Data: Demands $\{(d_i, F_i)\}_{i \in N}$, Capacities $\{c_f\}_{f \in M}$
Result: The Leximin-Optimal Allocation A
For $k \in \{1, \ldots, n\}$, let u^k denote the k^{th} lowest utility under an allocation;
for $k = 1$ to n **do**
  $\bar{u}^k \leftarrow$ Max u^k subject to $u^j = \bar{u}^j$ for all $j < k$;
end
return an allocation where $u^k = \bar{u}^k$ for all $k \in \{1, \ldots, n\}$;

In a sense, the leximin mechanism is an extension of the egalitarian equivalence principle put forward by Pazner and Schmeidler [166], in which one attempts to equalize all agent utilities (and maximize this utility value). This is what the leximin mechanism attempts in its first step of maximizing the minimum utility. However, sometimes the solution obtained is not Pareto optimal. The subsequent steps amend this solution to make it Pareto optimal, and eliminate any waste of resources. Without loss of generality, assume that the leximin mechanism chooses a non-wasteful allocation, i.e., under every deterministic assignment in its support agent i either receives d_i units from a facility in F_i or does not receive any units. Let us illustrate how the leximin mechanism works through an example.

Example 3.2. Suppose there are two facilities a and b with capacities $c_a = 1$ and $c_b = 2$, respectively, and four agents with demands $d_1 = 1$, $d_2 = 1$, $d_3 = 2$, and $d_4 = 1$.

Suppose agent 1 only accepts facility a ($F_1 = \{a\}$), agents 2 and 3 only accept facility b ($F_2 = F_3 = \{b\}$), and agent 4 accepts both facilities ($F_4 = \{a, b\}$).

It is clear that the minimum utility cannot be greater than $1/2$ because agents 2 and 3 must be assigned to facility b separately. Further, the randomized allocation $1/2 \cdot (a : \{1\}, b : \{3\}) + 1/2 \cdot (a : \{4\}, b : \{2\})$ (meaning that with probability $1/2$, agents 1 and 2 are assigned to facilities a and b, respectively, and with the remaining probability agents 4 and 3 are assigned to facilities a and b, respectively) gives utility $1/2$ to all agents. However, this is not sufficient for the allocation to be leximin-optimal. For instance, the allocation $1/2 \cdot (a : \{1\}, b : \{2, 4\}) + 1/2 \cdot (a : \{4\}, b : \{3\})$ increases the utility of agent 4 to 1 while keeping the utilities of the other agents at $1/2$, and is therefore better than the previous allocation in a lexicographic comparison of the sorted utility vector. While this new allocation is Pareto optimal, it is still not the leximin allocation. The leximin allocation in this example is $1/2 \cdot (a : \{1\}, b : \{2, 4\}) + 1/4 \cdot (a : \{1\}, b : \{3\}) + 1/4 \cdot (a : \{4\}, b : \{3\})$, which gives utility $1/2$ to agents 2 and 3, and utility $3/4$ to agents 1 and 4. Note that it achieves the same lowest and 2nd lowest utilities as the previous two allocations, but a greater 3rd lowest utility than both previous allocations.

3.3.1 Properties of The Leximin Mechanism

Bogomolnaia and Moulin [30] show that the leximin mechanism satisfies all four desiderata proposed above in their classic setting with one-to-one matchings, and unit demands and capacities. We now show that these properties continue to hold in our setting with many-to-one matchings, and arbitrary demands and capacities. In fact, in Section 3.3.2 we argue that they hold in an even more general setting.

Theorem 3.1. *The leximin mechanism satisfies proportionality, envy-freeness, Pareto optimality, and group strategyproofness.*

Proof. We first formally establish an intuitive property of leximin allocations: The allocation received by agent i is valued the most under the preferences of agent i compared to any other possible preferences. Formally, we show the following.

Lemma 3.1. *Let A denote the allocation returned by the leximin mechanism. Then for utility function u induced by any preferences, we have $u_i(A_i) \geqslant u(A_i)$.*

Proof. First, let A be deterministic. If $A_i \neq 0$, then due to the non-wastefulness of the leximin allocation, we must have $u_i(A_i) = 1 \geqslant u(A_i)$ for any utility function u. On the other hand, $A_i = 0$ implies $u_i(A_i) = u(A_i) = 0$ for all utility functions u. Hence, the lemma holds for all deterministic allocations. For randomized allocations, taking expectation on both sides yields that the lemma still holds. ■ (Proof of Lemma 3.1)

Proportionality. Consider the mechanism that allocates all available units to each agent with probability $1/n$, which gives each agent utility $1/n$.[4] Since the leximin mechanism maximizes the minimum utility that any agent receives, it must also give each agent at least $1/n$ utility. Hence, the leximin mechanism is proportional.

[4]This is because we assumed that the demand of every agent can be satisfied given all available units.

Envy-freeness. Suppose for contradiction that under an allocation A returned by the leximin mechanism, agent i envies agent j. That is, $u_i(A_j) > u_i(A_i)$. Now, Lemma 3.1 implies $u_j(A_j) \geqslant u_i(A_j) > u_i(A_i) \geqslant 0$. Let $0 < \varepsilon < (u_j(A_j) - u_i(A_i))/u_j(A_j)$.

Construct another allocation A' such that $A'_k = A_k$ for all $k \in N \setminus \{i, j\}$, $A'_i = A_j$, and $A'_j = 0$. Since agent i envied agent j, we have $d_i \leqslant d_j$, implying that A' is feasible. Note that agent i now has higher utility because $u_i(A'_i) = u_i(A_j) > u_i(A_i)$.

Construct an allocation A'' that realizes A with probability $1 - \varepsilon$ and A' with probability ε. Due to our construction of A'', we have that for every agent $k \in N \setminus \{i, j\}$, $u_k(A'') = u_k(A') = u_k(A)$. Further, for agent i we have $u_i(A''_i) > u_i(A_i)$. Also, for agent j we have

$$u_j(A''_j) = (1 - \varepsilon)u_j(A_j) > u_i(A_i).$$

Hence, switching from A to A'' preserves the utility achieved by every agent except agents i and j, and both agents i and j receive utility strictly greater than $u_i(A_i) = \min(u_i(A_i), u_j(A_j))$. That is, allocation A'' is strictly better than allocation A in the leximin ordering, which contradicts the leximin-optimality of A.

Pareto optimality. This follows trivially from the definition of leximin-optimality. Note that increasing the utility of an agent i without decreasing the utility of any other agent would improve the allocation in the leximin ordering. Since the allocation returned by the leximin mechanism is already leximin-optimal, it does not admit any Pareto improvements. Hence, the leximin mechanism is Pareto optimal.

Group Strategyproofness. This is the most non-trivial property to establish among the four desired properties. Under the true reports $(d_k, F_k)_{k \in N}$, let A denote the allocation returned by the leximin mechanism. Suppose a subset of agents $S \subseteq N$, whom we call manipulators, report false preferences $(d'_i, F'_i)_{i \in S}$; let $(u'_i)_{i \in S}$ denote the utility functions induced by the false preferences of the manipulators. Let A' denote the allocation returned by the leximin mechanism when agents in S misreport. Suppose for contradiction that every agent in S is strictly better off (under their true utility functions) by misreporting, i.e., $u_i(A'_i) > u_i(A_i)$ for every $i \in S$. Now, Lemma 3.1 implies that $u'_i(A'_i) \geqslant u_i(A'_i)$; thus, we have $u'_i(A'_i) > u_i(A_i)$ for every $i \in S$.

Before we derive a contradiction, we first observe that the leximin-optimality of an allocation implies Pareto optimality of any prefix of its sorted utility vector. Let $\text{pref}_A(i) = \{j \in N \mid u_j(A_j) \leqslant u_i(A_i)\}$ denote the prefix of agent i in allocation A.

Lemma 3.2 (Prefix Optimality)**.** *For an allocation X returned by the leximin mechanism and an agent $i \in N$, there does not exist an allocation X' such that some agent in $\text{pref}_X(i)$ is strictly better off under X' and no agent in $\text{pref}_X(i)$ is worse off.*

Proof. Assume without loss of generality that $u_i(X_i) < \max_{j \in N} u_j(X_j)$, otherwise the statement coincides with Pareto optimality. Suppose for contradiction that an allocation X' as in the statement of the lemma exists. Choose ε such that

$$0 < \varepsilon < \frac{1 - u_i(X_i)}{\min\{u_j(X_j) \mid u_j(X_j) > u_i(X_i)\}}.$$

Consider the allocation $X'' = (1-\varepsilon)\cdot X + \varepsilon \cdot X'$. Due to our choice of ε, we can see that for every agent $j \notin \mathrm{pref}_X(i)$, we have $u_j(X''_j) \geqslant (1-\varepsilon)u_j(X_j) > u_i(X_i)$. Further, we have $u_j(X''_j) \geqslant u_j(X_j)$ for every agent $j \in \mathrm{pref}_X(i)$ and $u_j(X''_j) > u_j(X_j)$ for some $j \in \mathrm{pref}_X(i)$.

We now show that X'' is strictly better than X in the leximin ordering. Choose agent $j^* \in \arg\min_{j \in \mathrm{pref}_X(i): u_j(X''_j) > u_j(X_j)} u_j(X_j)$. Break ties by choosing an agent with the smallest value of $u_j(X''_j)$, and if there are still ties, break them arbitrarily. Let $t = |\{k \in \mathrm{pref}_X(i) | u_k(X_k) < u_{j^*}(X_{j^*})\}| + |\{k \in \mathrm{pref}_X(i) | u_k(X''_k) = u_k(X_k) = u_j(X_j)\}|$. Then, one can check that allocations X and X'' match in the t lowest utilities, and allocation X'' has a strictly greater $(t+1)^{st}$ lowest utility. Thus, X'' is strictly better than X in the leximin ordering, which contradicts leximin-optimality of X. ■ (Proof of Lemma 3.2)

Fix a manipulator $i \in S$ that minimizes $u_i(A_i)$ among all $i \in S$ (break ties arbitrarily). Let us look at the set of all agents that are strictly better off under A' compared to A, and among these agents, choose an agent j that minimizes $u_j(A_j)$ (again, break ties arbitrarily). Now, agent i is also strictly better off under A'. Hence, by the definition of agent j, we have $u_j(A_j) \leqslant u_i(A_i)$. Since agent j is strictly better off under A', by prefix optimality of A (Lemma 3.2) we know there must exist an agent in $\mathrm{pref}_A(j)$ that is strictly worse off under A'. Among all agents in $\mathrm{pref}_A(j)$ that are worse off under A', choose an agent k that minimizes $u_k(A'_k)$ (again, break ties arbitrarily).

Now, we derive our contradiction by showing that prefix optimality of A' is violated. More precisely, we know that agent k is strictly better off under A compared to A'. We show that no agent in $\mathrm{pref}_{A'}(k)$ is worse off under A compared to A'.

First, note that for any manipulator $l \in S$, we have $u_{l'}(A'_l) \geqslant u_l(A'_l) > u_l(A_l) \geqslant u_i(A_i) \geqslant u_j(A_j) \geqslant u_k(A_k) > u_k(A'_k)$, where the third, fourth, and fifth transitions hold due to our choice of agents i, j, and k, respectively. Thus, no manipulator belongs to $\mathrm{pref}_{A'}(k)$. In other words, for every agent $l \in \mathrm{pref}_{A'}(k)$ we can denote its utility function (which is common between A and A') by u_l. Take an agent $l \in \mathrm{pref}_{A'}(k)$. If $u_l(A_l) < u_l(A'_l)$, then we have $u_l(A_l) < u_l(A'_l) \leqslant u_k(A'_k) < u_k(A_k) \leqslant u_j(A_j)$. Thus, agent l satisfies $u_l(A_l) < u_j(A_j)$, and is still better off under A' compared to A, which contradicts our choice of agent j. Therefore, $u_l(A_l) \geqslant u_l(A'_l)$ for every $l \in \mathrm{pref}_{A'}(k)$, and $u_k(A_k) > u_k(A'_k)$, contradicting prefix optimality of A'. ■ (Proof of Theorem 3.1)

While group strategyproofness is a strong game-theoretic requirement, an even stronger requirement has been studied in the literature. Under this stronger requirement, a group of manipulators should not be able to report false preferences that would lead to all manipulators being weakly happier and at least one manipulator being strictly happier. Bogomolnaia and Moulin [30] show that in the classical random assignment setting under dichotomous preferences, the leximin mechanism is group strategyproof according to this stronger requirement.[5] Unfortunately, the following example shows that this does not hold in our more general setting.

[5]While the strategyproofness result of Bogomolnaia and Moulin [30] more generally applies to strategic manipulations from both sides of the market, this is captured by our generalized results in Section 3.3.2.

Example 3.3. Suppose there are 9 agents with demands

$$(d_1, d_2, d_3, d_4, d_5, d_6, d_7, d_8, d_9) = (2, 4, 4, 4, 2, 2, 2, 1, 1),$$

and 3 facilities with capacities $(c_1, c_2, c_3) = (4, 2, 1)$. Let the dichotomous preferences of the agents be as follows: $F_i = \{1\}$ for $i \in \{1, 2, 3, 4\}$, $F_5 = \{1, 2\}$, $F_6 = F_7 = \{2\}$, $F_8 = \{2, 3\}$, and $F_9 = \{3\}$.

In this case, it can be checked that under the leximin allocation, the utilities of the agents are as follows: $u_i = 1/4$ for $i \in \{1, 2, 3, 4\}$, $u_5 = u_6 = u_7 = 5/12$, and $u_8 = u_9 = 1/2$.

Suppose agent 1 manipulates, and increases its demand to $d'_1 = 3$ units. Then, it can be checked that under the new leximin allocation, the utility of agent 1 through 4 remains $1/4$, the utility of agents 5 through 7 drops to $1/3$, and the utility of agents 8 and 9 increases to $5/8$. Thus, agent 1 and agent 9 form a successful group manipulation in which no agent is worse off, but agent 9 is strictly better off.

Similarly, Bogomolnaia and Moulin [30] also show that a leximin-optimal allocation always *Lorenz-dominates* any other allocation in their classic setting. Let us first define Lorenz dominance among allocations.

Lorenz dominance. For $k \in \{1, \ldots, n\}$, let u^k and v^k denote the k^{th} lowest utility in allocations A and B, respectively. We say that allocation A (weakly) Lorenz-dominates allocation B if $\sum_{i=1}^{k} u^i \geqslant \sum_{i=1}^{k} v^i$ for $k \in \{1, \ldots, n\}$.

We now show that in our setting there may not exist an allocation that weakly Lorenz-dominates every other allocation.

Example 3.4. Suppose there is a single facility with 3 available units, and there are four agents — namely, agents 1 through 4 — such that agent 1 demands all 3 units from the facility, while the remaining agents demand a single unit each. Suppose there exists an allocation A that weakly Lorenz-dominates every other feasible allocation. Then, in particular, it must achieve the maximum possible lowest utility. Hence, allocation A must assign agent 1 to the facility with probability 0.5, and assign the remaining agents to the facility simultaneously with the remaining probability 0.5. Thus, the sum of first three lowest utilities under A is 1.5. However, for the allocation that assigns agents 2 through 4 to the facility with probability 1, the sum of the three lowest utilities is 2, violating our assumption that A weakly Lorenz-dominates every other feasible allocation. Thus, in this case there does not exist any allocation that Lorenz-dominates every other allocation.

In general, the leximin allocation may not be unique, but all leximin allocations are equivalent in a sense formalized in the next result.

Theorem 3.2. *The utility of an agent is identical under all leximin allocations.*

Proof. Suppose for contradiction that there exist leximin-optimal allocations A and B such that the utilities of some agents do not match in the two allocations. Choose agent $i \in \arg\min_{i \in N : u_i(A_i) \neq u_i(B_i)} u_i(A_i)$, and break ties by choosing an agent with the smallest $u_i(B_i)$ (further ties can be broken arbitrarily). First, prefix optimality of A (Lemma 3.2) implies that agent i must be worse off under B, i.e., $u_i(B_i) < u_i(A_i)$. This is because otherwise there would exist an agent $j \in \text{pref}_A(i)$ that is strictly worse off under B. Agent j would satisfy $u_j(B_j) < u_j(A_j) \leqslant u_i(A_i) < u_i(B_i)$, and thus contradict our choice of agent i. Hence, we have $u_i(B_i) < u_i(A_i)$.

Now, consider the prefix of agent i in B, i.e., $\text{pref}_B(i)$. For every agent $j \in \text{pref}_B(i)$, either agent j has identical utility under A and B (i.e., $u_j(A_j) = u_j(B_j)$), or its utility changes in which case we must have $u_j(A_j) \geqslant u_i(A_i) > u_i(B_i) \geqslant u_j(B_j)$, where the first transition holds due to our choice of agent i. Hence, no agent in $\text{pref}_B(i)$ is worse off under A compared to B, and agent i is strictly better off under A compared to B. This violates prefix-optimality of B, which is a contradiction. Hence, the utility of each agent must be identical under all leximin-optimal allocations. ■

Crucially, this also implies that all leximin-optimal allocations satisfy equal number of agents in expectation, and allocate equal number of units in expectation.

3.3.2 A General Framework for Leximin

Theorem 3.1 established that the leximin mechanism satisfies four compelling desiderata in our classroom allocation setting. We observe that the proof of Theorem 3.1 only uses four characteristics of the classroom allocation setting (which are listed below). That is, the leximin mechanism (Algorithm 1) satisfies proportionality, envy-freeness, Pareto optimality, and group strategyproofness in all domains of fair division and mechanism design without money — with divisible or indivisible (or both types of) resources, and with deterministic or randomized allocations — that satisfy these four requirements.

We briefly describe a general framework in which our result holds. Let N denote the set of agents. There is a set of resources X, which may contain divisible resources, indivisible resources, or both. An allocation A assigns a *disjoint subset* of resources A_i to each agent i.[6] Denote the set of all feasible allocations by $\mathcal{A}$. Note that the use of randomized allocations may or may not be permitted in the domain; it does not affect our result. There is a set $\mathcal{P}$ of possible preferences that the agents may have over subsets of resources. Fix a mapping from each preference $P \in P$ to a utility function u_P consistent with P, and let $\mathcal{U} = \{u_P | P \in \mathcal{P}\}$ denote the corresponding set of possible utility functions. Then, our four requirements can be formalized as follows.

1. *Convexity.* The space of feasible allocations must be convex. That is, given two allocations $A, A' \in \mathcal{A}$, and $0 \leqslant \lambda \leqslant 1$, it should be possible to construct another feasible allocation $A'' \in \mathcal{A}$ such that $u_i(A''_i) = \lambda \cdot u_i(A_i) + (1 - \lambda)u_i(A'_i)$ for all agents $i \in N$. This typically holds if randomized allocations are allowed, or if resources are divisible.
2. *Equality.* The maximum utility achievable by each agent must be identical. Thus, for two agents $i, j \in N$, we require $\max_{A \in \mathcal{A}} u_i(A_i) = \max_{A \in \mathcal{A}} u_j(A_j)$. This prop-

[6]Obviously, only divisible resources can be split among multiple agents.

erty is required for proportionality, and is usually taken care of when translating the ordinal preferences of agents into cardinal utility functions.

3. *Shifting Allocations.* Given a feasible allocation $A \in \mathcal{A}$ and agents $i, j \in N$, it should be possible to construct another feasible allocation $A' \in \mathcal{A}$ where we take the resources allocated to agent j, and allocate them to agent i. That is, we must have $u_k(A'_k) = u_k(A_k)$ for all agents $k \in N \setminus \{i, j\}$, and $u_i(A'_i) \geqslant u_i(A_j)$. This property is required for envy-freeness.
4. *Optimal utilization.* Under a non-wasteful allocation $A \in \mathcal{A}$, an agent must derive the maximum possible utility from the allocation it receives. That is, we require $u_i(A_i) \geqslant u(A_i)$ for all possible utility functions $u \in \mathcal{U}$. Lemma 3.1 proves that this is satisfied in the classroom allocation setting. This assumption is perhaps the strongest, and is required for both envy-freeness and group strategyproofness.

Many papers study the leximin mechanism and establish (at least a subset of) the properties listed in Theorem 3.1 in a variety of domains, including resource allocation [28, 30, 99, 101, 133, 165, 192, 199, 200], cake cutting [56], and kidney exchange [186]. It can be checked that these domains satisfy our four requirements, and hence *the foregoing framework captures results from all of these papers.*

In addition, any general dichotomous preference setting — where each agent "accepts" a subset of feasible allocations for which it has utility 1, and "rejects" the rest for which it has utility 0 — is also captured under our general framework; and when agents have ordinal preferences over allocations, we only need to establish one translation to consistent cardinal utilities that satisfies the four requirements above.

Below, we briefly describe one special case of the general framework: fair resource allocation under Leontief preferences [99, 165], which is the focus of Chapter 4. Suppose there are m divisible resources, and each agent i demands them in fixed proportions given by a (normalized) demand vector $\boldsymbol{d} = (d_{i,1}, \ldots, d_{i,m})$ where $\max_{r \in \{1,\ldots,m\}} d_{i,r} = 1$. Thus, given an allocation $A_i = (A_{i,1}, \ldots, A_{i,m})$ (where $A_{i,r} \in [0, 1]$ denotes the fraction of resource r allocated to agent i), the utility to agent i is given by $u_i(A_i) = \min_{r \in \{1,\ldots,m\}} A_{i,r}/d_{i,r}$. To see that our four requirements are met, note that the space of feasible allocations is convex due to divisibility of resources, every agent can achieve a maximum utility of 1, and shifting allocations is permitted. Finally, a non-wasteful allocation always allocates resources in the demanded proportion. Thus, the utility to agent i is simply $A_{i,r}/d_{i,r}$ (which is identical for all r). Under any other normalized demand vector $d' = (d'_1, \ldots, d'_m)$ with $d'_{r^*} = 1$, the utility achieved would be at most $A_{i,r^*} \leqslant A_{i,r^*}/d_{i,r^*}$. Hence, the requirement of optimal utilization also holds.

Ghodsi et al. [99] prove that the leximin mechanism satisfies proportionality, envy-freeness, Pareto optimality, and strategyproofness in the foregoing setting, and in Chapter 4 (which is based on work that preceded the work presented in this chapter), we establish group strategyproofness. These results now directly follow from Theorem 3.1. Further, in Chapter 4, we study the variant where agents only derive utility for integral multiples of their required resource bundle, and show that no deterministic mechanism satisfies all four desiderata. Indeed, in our framework the convexity requirement is violated for deterministic allocations, but it is satisfied for randomized allocations. Hence, the randomized leximin mechanism would still satisfy all four desiderata.

3.4 Quantitative Efficiency of the Leximin Allocation

Theorem 3.1 establishes the leximin mechanism as a compelling solution, which simultaneously guarantees fairness, efficiency, and truthfulness. The fairness (proportionality and envy-freeness) and truthfulness guarantees are strong. But the notion of Pareto optimality is a relatively weak, qualitative notion of efficiency.

In our setting, there are two natural *quantitative* metrics of efficiency: the (expected) number of agents whose demands are met, and the (expected) number of total units allocated. Optimizing the former metric is clearly desirable as it represents the social welfare achieved by the mechanism. The latter metric is important when the units being allocated are valuable and scarce (this is clearly the case when the units in question are classrooms). Furthermore, in the classroom allocation setting, the number of units allocated is proportional to the number of students served.

Indeed, in our setting it is not unnatural to consider directly optimizing either metric. In particular, such an optimization would always lead to a Pareto optimal allocation. However, it is easy to observe that directly optimizing either metric fails to achieve one or more of our four desired properties. Recall Example 3.1, which already showed that maximizing the number of allocated units violates proportionality, envy-freeness, and strategyproofness; the next example deals with the other metric.

Example 3.5 (Maximizing the number of satisfied agents)**.** Suppose there is a single facility with 2 available units, and there are four agents, namely, agents 1 through 4. Agents 1 through 3 each demand a single unit from the facility, while agent 4 demands both units. In order to maximize the number of satisfied agents we must allocate a single unit to two of the agents in $\{1, 2, 3\}$, while leaving agent 4 unallocated. It is easy to see that both proportionality (with respect to agent 4) and envy-freeness (with respect to the unallocated agent in $\{1, 2, 3\}$) are violated.

In the above example, proportionality is clearly violated, but it seems that the violation of envy-freeness is the result of tie-breaking. Indeed, as previously mentioned, the utilitarian mechanism [31, 94] that uniformly randomizes over all deterministic allocations maximizing the number of satisfied agents achieves envy-freeness along with strategyproofness. We note that strategyproofness would also hold if ties were broken according to a lexicographic order over the agents. Here, we provide a short proof of these results for curious readers.

Observation 3.1. *The mechanism that returns an allocation maximizing the number of satisfied agents and breaks ties according to a lexicographic preference over agents is strategyproof. Breaking ties uniformly at random preserves strategyproofness, and simultaneously achieves envy-freeness.*

Proof. Let A denote the allocation returned by the mechanism under consideration with lexicographic tie-breaking.

Strategyproofness. Suppose agent $i \in N$ is not satisfied under A. Suppose agent i manipulates, which results in allocation A' satisfying agent i. Let k and k' denote the number of agents satisfied in A and A', respectively. Since agent i cannot decrease its demanded number of units, any subset of agents satisfiable after the manipulation is also satisfiable before the manipulation. Hence, $k \geqslant k'$. However, allocation A does not assign agent i to any facility, and therefore must be feasible after the manipulation. Thus, $k' \geqslant k$, implying $k = k'$. Finally, note that the subset of agents satisfied by A' was feasible before manipulation, but was not chosen because the subset of agents satisfied under A was better in the lexicographic preference. Since A is a feasible allocation after manipulation, it would still be preferred to A' under the same lexicographic preference, thus establishing a contradiction.

Suppose the mechanism returns an allocation A that uniformly randomizes over all allocations maximizing the number of satisfied agents. Let A' denote the corresponding (uniformly randomizing) allocation when agent i manipulates. If agent i is satisfied with probability 1 under A, then it has no incentive to manipulate. Otherwise, there exists an allocation in the support of A that does not satisfy agent i. Observing that this allocation is feasible after manipulation, and that every subset of agents satisfiable after manipulation is also satisfiable before manipulation, we again get $k = k'$. Moreover, since agent i cannot decrease its demand, the number of allocations in the support of A' in which agent i is satisfied is at most the number of such allocations in A. Since both A and A' uniformly randomize over allocations in their support, it is clear that agent i cannot increase its utility by manipulating.

Envy-freeness. Consider agents $i, j \in N$. Suppose for contradiction that agent i envies agent j. Let I denote the set of deterministic allocations in the support of A in which agent i is assigned to a facility, while agent j is unassigned. Let J denote the set of deterministic allocations in the support of A in which agent j is assigned to a facility *that is acceptable to agent i*, while agent i is unassigned. Let p_I and p_J denote the probabilities by which A executes an assignment from I and J, respectively. Then, we must have $p_J > p_I$. However, since agent i envies agent j, we must have $d_j \geqslant d_i$. Thus, taking an allocation from J, and replacing agent j with agent i must form a feasible allocation. Thus, $|I| \geqslant |J|$. Due to uniform randomization over all allocations in the support, we get $p_I \geqslant p_J$, which is a contradiction. ■ (Proof of Observation 3.1)

While the utilitarian mechanism seems intriguing, recall that in Example 3.5 the demand of agent 4 was met with zero probability, suggesting that the mechanism is biased against agents with larger demands. Bogomolnaia et al. call this effect the "tyranny of the majority". While such a bias may be acceptable in some settings, in other settings — classroom allocation, in particular — it is problematic. The bias is formally captured by noting that the utilitarian mechanism *violates proportionality*.

The discussion above leads us to a natural question: *How well does the leximin mechanism perform with respect to the two quantitative notions of efficiency, namely the number of satisfied agents and the number of allocated units?* We are interested in the worst case over problem instances, but since the leximin mechanism is randomized, we can consider the performance under the worst deterministic allocation in the support of the randomized leximin allocation, and the performance in expectation. Unsurprisingly, the worst allocation in the support can be simultaneously bad in terms of both metrics; in the following example, both metrics achieve arbitrarily low fractions of their respective optimums.

Example 3.6 (Efficiency of allocations in the support of the leximin allocation)**.** Suppose there are $k+4$ agents and two facilities. The capacities of the two facilities are $c_1 = k$ and $c_2 = k^2$. The preferences of the agents are as follows.

$$(d_i, F_i) = \begin{cases} (1, \{1\}) & \text{if } i \in \{1, \ldots, k\}, \\ (k, \{1\}) & \text{if } i = k+1 \text{ or } k+2, \\ (1, \{2\}) & \text{if } i = k+3, \\ (k^2, \{2\}) & \text{if } i = k+4. \end{cases}$$

Clearly, a maximum of $k+1$ agents can be satisfied, and a maximum of $k+k^2$ units can be allocated. It is easy to check that under the leximin allocation, agents 1 through $k+2$ should be assigned to facility 1 with probability $1/3$ each, while agents $k+3$ and $k+4$ should be assigned to facility 2 with probability $1/2$ each. However, this implies that the support of the leximin allocation must include a deterministic allocation in which agent $k+3$ is assigned to facility 2 while one of agents $k+1$ and $k+2$ is assigned to facility 1 (and the remaining agents are unassigned). In this allocation, the number of agents satisfied is a mere $2/(k+1)$ fraction of the optimum, and the number of units allocated is also a mere $(k+1)/(k+k^2) = 1/k$ fraction of the optimum. Thus, both approximation ratios converge to 0 as k goes to infinity.

Let us therefore consider the worst-case (over instances) performance of the leximin mechanism in expectation (over the randomness of the mechanism). We can show that approximating (in expectation) the maximum number of satisfied agents is directly at odds with proportionality — recall that this is exactly the property that the utilitarian mechanism [31, 94] fails to achieve.

Example 3.7 (Proportionality and maximizing the number of satisfied agents)**.** Suppose there is a single facility with k units available, and there are $k+k^2$ agents, k of which require 1 unit each while the other k^2 agents require all k units each. Any proportional mechanism must allocate the k units to each of the k^2 agents demanding them with probability at least $1/(k+k^2)$. Hence, such a mechanism satisfies a single agent with probability at least $k^2/(k+k^2)$, and at most k agents with the remaining probability. Therefore, the expected number of satisfied agents is at most $k^2/(k+k^2) + k \cdot k/(k+k^2) \leqslant 2$. However, a maximum of k agents could be satisfied simultaneously. Hence, any proportional mechanism (including the leximin mechanism) achieves an approximation ratio of at most $2/k$ for the number of satisfied agents. This ratio goes to 0 as k goes to infinity.

In contrast, we make the following conjecture for the expected number of *units* allocated by the leximin mechanism:

Conjecture 3.1. *The expected number of units allocated by the leximin mechanism* 2*-approximates the maximum number of units that can be allocated simultaneously by any non-wasteful allocation (in the worst case over instances).*

The conjecture is based on millions of randomly generated instances. In all of these instances, the leximin mechanism allocated, in expectation, at least half of the optimal number of units. While the conjecture is still open, we are able to prove a slightly weaker 4-approximation result.

Theorem 3.3. *The expected number of units allocated by the leximin mechanism* 4*-approximates the maximum number of units that can be allocated simultaneously by any non-wasteful allocation (in the worst case over instances).*

Proof. Let us first prove a 2-approximation in the case of a single facility to gain some intuition. Let c denote the capacity of the facility, and D denote the maximum number of units allocated by a non-wasteful allocation. If all the deterministic assignments in the support of the leximin allocation allocate at least $D/2$ units, then the result follows trivially. Suppose a deterministic assignment allocates $t < D/2 \leqslant c/2$ units to agents in $S \subseteq N$, and is realized with probability p. Hence, it is clear that $N \setminus S \neq \emptyset$. Due to Pareto optimality of the leximin allocation, an allocation that does not assign any agent in $N \setminus S$ to the facility must assign all agents in S to the facility. That is, there is a unique such allocation, which is realized with probability p. Further, due to the nature of the leximin allocation, every agent in $N \setminus S$ must also be assigned to the facility with probability at least p, implying that $p \leqslant 1/2$. Thus, with probability $p \leqslant 1/2$ the mechanism allocates t units, and with the remaining probability $1 - p$ the mechanism assigns at least one agent in $N \setminus S$ to the facility, thus allocating more than $c - t$ units. Hence, the expected number of units allocated is at least $t \cdot 1/2 + (c - t) \cdot 1/2 = c/2 \geqslant D/2$.

However, generalizing this proof to achieve a "per facility" constant approximation is difficult. Instead, our proof below works in three steps.

1. We fix an arbitrary (deterministic) allocation A^* that maximizes the number of units allocated.
2. Next, after adding certain "virtual allocated units" to each site (derived based on A^*), the expected number of units allocated by the leximin mechanism 2-approximates the number of units allocated under A^* *on each facility individually*.
3. Finally, we show that the expected number of virtual units added overall is no more than the expected number of units allocated by the leximin mechanism, thus establishing the 4-approximation result.

Let A^* denote an arbitrary deterministic allocation that maximizes the number of units allocated. For a facility $f \in M$, let $Z(f) = \{i \in N | A^*_i = f\}$ denote the set of agents assigned to facility f under A^*. Let L denote the leximin allocation, which executes deterministic allocation L^k with probability p_k for $k \in \{1, \ldots, T\}$. We are now ready for our main lemma. For a facility $f \in M$, the number of "virtual units" we add is the expected number of units allocated by the leximin mechanism to the agents in $Z(f)$ (at any facility). We show that the expected number of units allocated by the leximin mechanism at facility f and the number of virtual units for facility f together 2-approximate the number of units allocated by A^* at facility f, for each $f \in M$.

Lemma 3.3. *For a facility $f \in M$ we have:*

$$\sum_{k=1}^{T} p_k \left(\sum_{i \in N : L_i^k = f} d_i + \sum_{i \in Z(f) : L_i^k \neq 0} d_i \right) \geqslant \frac{1}{2} \sum_{i \in Z(f)} d_i.$$

Proof. Let us consider two cases.

Case 1: For every $i \in Z(f)$, $d_i \leqslant c_f/2$. In this case we can show that

$$\sum_{i \in N : L_i^k = f} d_i + \sum_{i \in Z(f) : L_i^k \neq 0} d_i \geqslant \frac{1}{2} \sum_{i \in Z(f)} d_i \tag{3.1}$$

for each $k \in \{1, \ldots, T\}$. If $L_i^k \neq 0$ for every $i \in Z(f)$, then the second term in the LHS of Equation (3.1) is at least $\sum_{i \in Z(f)} d_i$. Otherwise, let $L^k(\gamma) = 0$ for some $\gamma \in Z(f)$. By the Pareto optimality of L^k, we know that the demand of agent γ must be greater than the number of unallocated units at facility f in L^k, i.e.,

$$d_\gamma > c_f - \sum_{i \in N : L_i^k = f} d_i.$$

Using $d_\gamma < c_f/2$, we get that the first term in the LHS of Equation (3.1) greater than the RHS. Hence, in either case Equation (3.1) holds.

Case 2: There exists an agent $\gamma \in Z(f)$ such that $d_\gamma > c_f/2$. Let us define two sets.

1. $I = \{k \in \{1, \ldots, T\} \mid L_\gamma^k \neq 0\}$
2. $J = \{k \in \{1, \ldots, T\} \mid L_\gamma^k = 0 \text{ and } \sum_{i \in N : L_i^k = f} d_i < c_f/2\}$

Furthermore, let $p_I = \sum_{k \in I} p_k$ and $p_J = \sum_{k \in J} p_k$. We claim that $p_I \geqslant p_J$. Note that p_I is precisely the probability that agent γ is satisfied under the leximin allocation.

Suppose for contradiction that $p_I < p_J$. Take some $\ell \in J$, and let $W = \{i \in N \mid L_i^\ell = f\}$. From the definition of J, we know that each agent $i \in W$ must satisfy $d_i < c_f/2$. Further, for each $k \in J$ facility f has more than $c_f/2$ units unallocated in L^k. Hence, by the Pareto optimality of the leximin allocation, every agent in W must be assigned to some facility in L^k for every $k \in J$. Importantly, this implies that every agent in W has probability at least $p_J > p_I$ of being assigned to a facility under the leximin allocation.

Now, fix a small $\varepsilon > 0$, and consider a new randomized allocation $\tilde{L}$ that executes deterministic allocations $L^1, \ldots, L^{k-1}, L^k, L^{k+1}, \ldots, L^T$, and L^{T+1} with probabilities $p_1, \ldots, p_{k-1}, (1-\varepsilon)p_k, p_{k+1}, \ldots, p_T$, and εp_k, respectively, where

$$L_i^{T+1} = \begin{cases} L_i^k & \text{if } L_i^k \neq f \\ 0 & \text{if } L_i^k = f \text{ and } i \neq \gamma \\ 1 & \text{otherwise.} \end{cases}$$

Note that f must be an acceptable site to agent γ because $\gamma \in Z(f)$. Hence, allocation L_i^{T+1} respects the preferences of the agents. It is easy to check that the capacity constraint at each facility (including facility f) is also respected. Essentially, we replace all the agents assigned at facility f in L^k by a single agent γ. For a sufficiently small $\varepsilon > 0$, we can see that:

1. Agent γ has a strictly higher probability of being assigned to a facility under $\tilde{L}$ than under L (under L, it is assigned to a facility with probability exactly p_I).
2. An agent $i \neq \gamma$ that is assigned to a facility with probability $p \leqslant p_I$ (thus, from the above argument $L_i^k \neq f$) has the same probability of being assigned to a facility under $\tilde{L}$ as under L.
3. All the remaining agents were assigned to a facility with probability strictly more than p_I under L, and their probabilities remain strictly greater than p_I under $\tilde{L}$.

However, this contradicts the fact that L is a leximin-optimal allocation. This is essentially a consequence of the prefix optimality of L (Lemma 3.2). Hence, we have $p_I \geqslant p_J$, as claimed.

With this claim in hand, we can show the required inequality. Let us consider the sum in the LHS.

$$\sum_{k=1}^{T} p_k \left(\sum_{i \in N : L_i^k = f} d_i + \sum_{i \in Z(f) : L_i^k \neq 0} d_i \right).$$

We break the summation over $k \in I$, $k \in J$, and $k \in \{1, \ldots, T\} \setminus (I \cup J)$. For each $k \in I$, we have $L_\gamma^k \neq 0$. Hence, the term inside the brackets is at least d_γ. For each $k \in J$, we have $L_\gamma^k = 0$. Hence, the term inside the brackets, which is no less than the number of units allocated at facility f in L^k, must be at least $c_f - d_\gamma$. Finally, from definitions of I and J, it follows that the term inside the brackets is at least $c_f/2$ for every $k \in \{1, \ldots, T\} \setminus (I \cup J)$. Hence, we have that the LHS is at least

$$\begin{aligned}
&\sum_{k \in I} p_k \cdot d_\gamma + \sum_{k \in J} p_k \cdot (c_f - d_\gamma) + \sum_{k \in \{1,\ldots,T\} \setminus (I \cup J)} p_k \cdot \frac{c_f}{2} \\
&\quad = p_I \cdot d_\gamma + p_J \cdot (c_f - d_\gamma) + (1 - p_I - p_J) \cdot \frac{c_f}{2} \\
&\quad = (p_I - p_J) \cdot d_\gamma + (1 - p_I + p_J) \cdot \frac{c_f}{2} \\
&\quad \geqslant (p_I - p_J) \cdot \frac{c_f}{2} + (1 - p_I + p_J) \cdot \frac{c_f}{2} = \frac{c_f}{2} \geqslant \frac{1}{2} \sum_{i \in Z(f)} d_i,
\end{aligned}$$

where the third transition holds because $p_I \geqslant p_J$ and $d_\gamma \geqslant c_f/2$. Thus, we have proved that the lemma holds in both the cases we considered. ■ (Proof of Lemma 3.3)

Lemma 3.3 holds for every facility individually. Summing over all facilities, we get:

$$\sum_{f \in M} \sum_{k=1}^{T} p_k \left(\sum_{i \in N : L_i^k = f} d_i + \sum_{i \in Z(f) : L_i^k \neq 0} d_i \right) \geqslant \frac{1}{2} \sum_{f \in M} \sum_{i \in Z(f)} d_i. \tag{3.2}$$

In Equation (3.2), we have

$$\begin{aligned}
\text{LHS} &= \sum_{k=1}^{T} p_k \cdot \left(\sum_{f \in M} \sum_{i \in N : L_i^k = f} d_i + \sum_{f \in M} \sum_{i \in Z(f) : L_i^k \neq 0} d_i \right) \\
&= \sum_{k=1}^{T} p_k \cdot \left(\sum_{i \in N : L_i^k \neq 0} d_i + \sum_{i \in N : A_i^* \neq 0, L_i^k \neq 0} d_i \right) \\
&\leqslant 2 \cdot \sum_{k=1}^{T} p_k \left(\sum_{i \in N : L_i^k \neq 0} d_i \right),
\end{aligned}$$

$$\text{RHS} = \frac{1}{2} \sum_{f \in M} \sum_{i \in Z(f)} d_i = \frac{1}{2} \sum_{i \in N : A_i^* \neq 0} d_i.$$

Note that LHS is at most twice the expected number of units allocated by the leximin mechanism, and RHS is half the number of units allocated by A^*. Hence, the expected number of units allocated by the leximin mechanism 4-approximates the maximum number of units allocated by a non-wasteful allocation. ■

While we strongly believe that the approximation ratio of Theorem 3.3 can be improved from 4 to 2, it can easily be seen that a proportional or envy-free mechanism (including the leximin mechanism) cannot achieve an approximation ratio better than 2. Consider the case of a single facility with $2k$ units, and $k+1$ agents, one of which requires all $2k$ units while the rest require $k+1$ units each. Clearly any proportional or envy-free mechanism must assign each agent demanding $k+1$ units alone to the facility with probability at least $1/(k+1)$. Hence, the expected number of allocated units cannot be more than $(k+1) \cdot k/(k+1) + 2k \cdot 1/(k+1) \leqslant k+2$, while a maximum of $2k$ units can be allocated simultaneously. This lower bound on the approximation ratio tends to 2 as k tends to infinity.

3.5 Complexity and Implementation

Recall that our classroom allocation setting is a generalization of the classic setting of random assignment under dichotomous preferences studied by Bogomolnaia and Moulin [30] (which can be viewed in our model as restricting agents to have unit demands and facilities to have unit capacities). In the classic setting, leximin allocations can be computed in polynomial time by leveraging the Birkhoff von-Neumann theorem [25, 197].

In contrast, an immediate reduction from PARTITION shows that computing the leximin allocation is $\mathcal{NP}$-hard in our generalized setting. Indeed, consider an instance of PARTITION: given a set S of n integers that sum to $2T$ for $T \in \mathbb{N}$, one needs to decide if there exists a subset $S' \subseteq S$ whose elements sum to T. Construct an instance of our problem in which a single facility has T available units and there are n agents whose demands correspond to the elements of S. Then, the leximin allocation would assign each agent to the facility with probability at least $1/2$ if and only if there exists a partition of S.

The standard approach to computing the leximin allocation (see, e.g., [153]) is to successively solve linear programs (LPs) in order to maximize the lowest utility, subject to that maximize the second lowest utility, and so on. While previous work focused on establishing polynomial running time of this approach in various domains, in our domain this task is $\mathcal{NP}$-complete. Hence, in the remainder of the section, we focus on designing optimized heuristics for computing the leximin allocation in the classroom allocation setting. We use a variable p_i to denote the probability that agent i is satisfied, for every $i \in N$. In a naïve implementation, we can include a variable x_A for every possible deterministic assignment $A \in \mathcal{A}$ that represents the probability of executing A, and write $p_i = \sum_{A \in \mathcal{A}: A_i \neq 0} x_A$. However, the number of feasible deterministic allocations can be roughly $(m+1)^n$, which makes the LPs extremely large even for moderately large values of m and n.

Crucially, note that we only care about whether a given agent is satisfied in a deterministic allocation, and not about the facility to which the agent is assigned. In other words, two deterministic allocations that satisfy identical subsets of agents are, in some sense, equivalent. This is due to the dichotomous nature of the preferences of agents over facilities. This observation leads us to our first algorithm, presented as Algorithm LEXIMINPRIMAL, which works as follows. First, we compute the collection of "feasible subsets" of agents, i.e., subsets of agents that can be satisfied simultaneously. Let $\mathcal{S} = \{S \subseteq N \mid \exists A \in \mathcal{A} \text{ s.t. } \forall i \in S, A_i \neq 0\}$. Checking feasibility of a given subset of agents S can be encoded as an integer linear program (ILP), presented as FEASIBILITYILP in the algorithm, which checks if agents in S can be assigned to one of their acceptable facilities while respecting the capacity constraints. Note that a feasible solution to FEASIBILITYILP also provides an assignment A_S that satisfies S.

Finally, we form an LP, which we call PRIMALLP, in which variable x_S denotes the probability by which $S \subseteq N$ is satisfied, and express the individual agent utilities as $p_i = \sum_{S \subseteq N: i \in S} x_S$ for $i \in N$. The algorithm maintains a set of agents R whose utilities in the leximin allocation it has not yet found, and stores the utility of each agent $i \in N \setminus R$ as p_i^*. In each iteration, the algorithm maximizes the (next) minimum utility of agents in R while keeping the utilities of agents in $N \setminus R$ intact, stores the utilities of agents that have the next minimum utility, and removes them from R.

The algorithm clearly terminates because any optimal solution to PRIMALLP must set $p_i = M$ for at least one $i \in R$. Hence, $|R|$ decreases by at least 1 in every iteration. Further, if M is the optimal objective value of PRIMALLP, then an observation from the convex optimization literature states that there must exist at least one $j \in R$ that has utility M in all optimal solutions to PRIMALLP, and in particular, in the actual leximin

ALGORITHM 2: LEXIMINPRIMAL

Data: Demands $\{(d_i, F_i)\}_{i \in N}$, Capacities $\{c_j\}_{j \in M}$
Result: The Leximin Allocation A
Solve FEASIBILITYILP for each $S \subseteq N$, and let $\mathcal{S} \leftarrow$ the set of maximal feasible subsets of N;
For each $S \in \mathcal{S}$, $A_S \leftarrow$ the assignment returned by FEASIBILITYILP on S;
$R = N$;
$p_i^* = 0, \forall i \in N$;
do
  $(M, \{p_i\}_{i \in R}, \{x_S\}_{S \in \mathcal{S}}) \leftarrow$ Strictly complementary solution to PRIMALLP in the box below;
  $p_i^* = M, \forall i \in R : p_i = M$;
  $R = R \setminus \{i \in N | p_i = M\}$;
  if $R = \emptyset$ **then**
    return the randomized allocation where A_S is executed with probability x_S for each $S \in \mathcal{S}$;
  end
while $R \neq \emptyset$;

PRIMALLP:
Maximize M subject to $p_i \geqslant M, \forall i \in R$ $p_i = p_i^*, \forall i \in N \setminus R$ $p_i = \sum_{S \in \mathcal{S}, i \in S} x_S, \forall i \in N$ $\sum_{S \in \mathcal{S}} x_S = 1$ $x_S \geqslant 0, \forall S \in \mathcal{S}$

FEASIBILITYILP:
$\sum_{f \in F_i} y_{i,f} \geqslant 1, \forall i \in S$ $\sum_{i \in S : f \in F_i} d_i \cdot y_{i,f} \leqslant c_f, \forall f \in M$ $y_{i,f} \in \{0, 1\}, \forall i \in S, f \in F_i$

allocation too.[7] Our use of a strictly complementary solution to PRIMALLP ensures that we have $p_j = M$ only if it holds in all optimal solutions.[8] Thus, Algorithm LEXIMINPRIMAL always makes "safe" choices, and correctly returns a leximin allocation. Finally, note that the values of p_i^* from one iteration are used to compute p_i^* in the next iteration. While this may lead to an exponential blowup in the length of their binary representation, it does not affect the running time of our algorithm due to a result by Tardos [193].[9] Interestingly, note that the choices made by the algorithm do not affect agent utilities in the returned leximin allocation due to Theorem 3.2.

[7]If for every $j \in R$ there exists a solution to PRIMALLP with $p_j > M$, a positive convex combination of these solutions would be a feasible solution with a strictly greater objective value, which is a contradiction.

[8]Strictly complementary solutions can be found by using any interior point method based on the central trajectory [95], by using a trick due to Freund et al. [96] which requires solving a single LP using any off-the-shelf solver, or by solving one LP for each $i \in R$ to check if p_i can be made greater than M in some optimal solution to PRIMALLP.

[9]This result shows that the running time of an interior point method is independent of the bit length of values on the right hand side of an LP, which is where the p_i^* are used in PRIMALLP.

We employ two further optimizations to reduce the running time of LEXIMINPRIMAL: i) solving FEASIBILITYILP on different subsets of agents in the decreasing order of their sizes, and only solving it for $S \subseteq N$ if none of its strict supersets are already found to be feasible, and ii) only using *maximal* feasible subsets in $\mathcal{S}$ because Pareto optimality prevents the leximin allocation from using any non-maximal subset.

ALGORITHM 3: LEXIMINDUAL

Data: Demands $\{(d_i, F_i)\}_{i \in N}$, Capacities $\{c_j\}_{j \in M}$
Result: The Leximin Allocation A
$R = N$;
$p_i^* = 0, \forall i \in N$;
do
  $(M, \{\alpha_i\}_{i \in R}) \leftarrow$ Strictly complementary solution to DUALLP in the box below;
  $p_i^* = M, \forall i \in R : \alpha_i > 0$;
  $R = R \setminus \{i \in N | p_i = M\}$;
while $R \neq \emptyset$;
$\widehat{\mathcal{S}} \leftarrow \{S \subseteq N |$ oracle of DUALLP was called on S in the last iteration of the loop$\}$;
For each $S \in \widehat{\mathcal{S}}$, $A_S \leftarrow$ the assignment returned by the oracle when it was called on S;
$\{x_S\}_{S \in \widehat{\mathcal{S}}} \leftarrow$ Solution to FINALLP in the box below;
return the randomized allocation where A_S is executed with probability x_S for each $S \in \widehat{\mathcal{S}}$;

DUALLP:	Oracle for DUALLP:	FINALLP:
Min. $M = \delta - \sum_{i \in N \setminus R} p_i^* \cdot \beta_i$ subject to $\sum_{i \in R} \alpha_i = 1$ $-\alpha_i - \gamma_i = 0, \forall i \in R$ $-\beta_i - \gamma_i = 0, \forall i \in N \setminus R$ $\delta + \sum_{i \in S} \gamma_i \geqslant 0, \forall S \in \mathcal{S}$ $\alpha_i \geqslant 0, \forall i \in R$	Max. $\sum_{i \in N} \gamma_i \cdot \left(\sum_{j \in F_i} y_{i,j} \right)$ subject to the constraints of FEASIBILITYILP	Find a feasible solution to $p_i = p_i^*, \forall i \in N$ $p_i = \sum_{S \in \widehat{\mathcal{S}}, i \in S} x_S, \forall i \in N$ $\sum_{S \in \widehat{\mathcal{S}}} x_S = 1$ $x_S \geqslant 0, \forall S \in \widehat{\mathcal{S}}$

Next, we present another algorithm that, instead of solving PRIMALLP, solves its dual. This is presented as Algorithm LEXIMINDUAL. Note that PRIMALLP has polynomially many constraints and exponentially many variables. Correspondingly, its dual (DUALLP) has polynomially many variables and exponentially many constraints (in particular, one constraint for each $S \in \mathcal{S}$). We can identify the tight primal constraints ($p_i = M$ for $i \in R$) by simply checking if the corresponding dual variable is strictly positive ($\alpha_i > 0$) due to the strict complementary slackness conditions. We solve DUALLP using the Ellipsoid algorithm [121], which makes polynomially many calls to an "oracle" for finding a violated constraint (if one exists) given any values of the variables. Crucially, we observe that finding $S \in \mathcal{S}$ that corresponds to the most violated constraint can be encoded as an ILP, presented along with the algorithm. We use $\widehat{\mathcal{S}}$ to denote the polynomial-size collection of subsets of agents on which the oracle is called by the Ellipsoid algorithm. There are three special advantages of the oracle:

1. Since the oracle includes feasibility constraints, we can avoid the initial (computationally expensive) stage of LEXIMINPRIMAL solving FEASIBILITYILP for 2^n subsets of agents, and instead solve only polynomially many ILPs for subsets in $\widehat{\mathcal{S}}$.
2. Since LEXIMINDUAL makes only polynomially many calls to the oracle, the overall space complexity is polynomial. In particular, the returned leximin allocation randomizes over *polynomially many subsets of agents* (i.e., it is *sparse*), making it more feasible to store and implement the allocation in practice.
3. In special cases such as the case of unit demands and capacities (i.e., the classic random assignment setting studied by Bogomolnaia and Moulin [30]), the oracle can be encoded as a polynomial-size LP by leveraging the Birkhoff von-Neumann theorem [25, 197], which would automatically make the overall running time of LEXIMINDUAL polynomial.

In the next section, we show that LEXIMINDUAL is actually drastically superior to LEXIMINPRIMAL in terms of running time.

3.6 Experiments

Our goal in this section is to empirically compare algorithms LEXIMINPRIMAL and LEXIMINDUAL, as well as evaluate the performance of the leximin allocation in terms of the number of satisfied agents and the number of allocated units.

In our experiments, we vary the number of agents n from 5 to 300.[10] Note that the largest school district in the US (by the number of charter schools) is the Los Angeles Unified School District (LAUSD) which has 241 charter schools.[11] We observe that in practice the number of facilities varies from about $5n$ (for LAUSD) to about $20n$ (for PUSD). Thus, we select m uniformly at random from the interval $[5n, 20n]$. Next, we fit Poisson distributions to the real-world demands and capacities data from PUSD, and use them to generate demands and capacities in our experiments. For the dichotomous preferences of agents over facilities, we observe that in the PUSD data certain facilities were inherently more desirable than others, and were accordingly accepted by many charter schools. We thus generate a "quality parameter" for each facility in $[0, 1]$ from the beta distribution with both parameters equal to 5, and have each agent accept the facilities (which have sufficient capacity to meet its demand) with probabilities proportional to their qualities. For each value of n, the values in all our graphs are averaged over 500 simulations. We use MATLAB to obtain strictly complementary solutions to linear programs, and CPLEX to solve integer linear programs. Our experiments are performed on an Intel PC with dual core, 3.10 GHz processors, and 8 GB RAM.

Figure 3.1 compares the running time of algorithms LEXIMINPRIMAL and LEXIMINDUAL. Note that the running time of LEXIMINPRIMAL increases extremely quickly as

[10] We use $n = 5, 10, 15$ for LEXIMINPRIMAL as it fails to run beyond that, and evaluate LEXIMINDUAL further on $n = 50, 100, 150, 200, 250, 300$.

[11] Refer to `leximin:http://goo.gl/Bu0pz9` and `leximin:http://goo.gl/ILJupc`

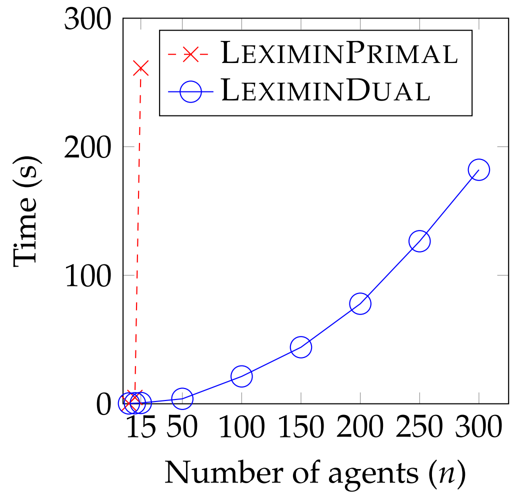

Figure 3.1: Running time of LEXIMINPRIMAL and LEXIMINDUAL.

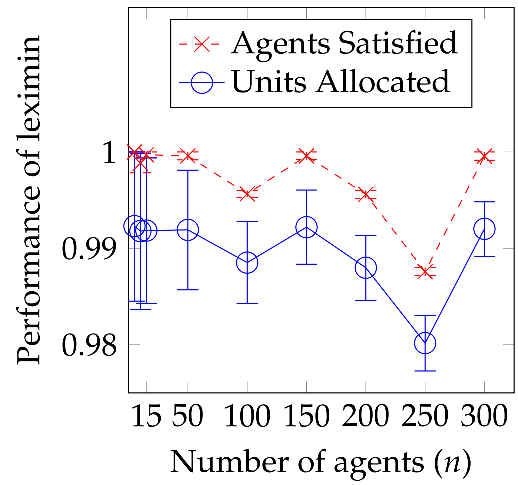

Figure 3.2: Performance of the leximin allocation as a fraction of the optimum.

n grows, making it infeasible to run the algorithm beyond $n = 15$. In contrast, LEXIMINDUAL solves instances with $n = 300$ (recall that this is larger than any real-world instance) in just a little over 3 minutes. This is a direct result of the fact that LEXIMINDUAL ends up solving less than 1% of the ILPs solved by LEXIMINPRIMAL, and solving ILPs is the bottleneck in both algorithms. Another interesting fact is that the number of times the loop in LEXIMINDUAL (or in LEXIMINPRIMAL) runs is equal to the number of distinct utility values in the leximin solution, because all agents with identical utilities are removed in a single iteration. The number of iterations required is less than 3 on average in our simulations. We also remark that even if the Proposition 39 process scaled to the state level, California has approximately 1130 charter schools overall,[11] and LEXIMINDUAL can also solve such huge instances in less than 2 hours (this result is averaged over 10 simulations).

Next, in Figure 3.2 we show the ratios of the expected number of agents satisfied and the expected number of units allocated by the leximin mechanism to the maximum possible values of the respective metrics. Remarkably, both ratios stay above a whopping 0.98 on average, which is significantly better than the upper bounds on the worst-case (over possible instances) performance of the leximin mechanism (almost 0 for the expected number of agents satisfied and 1/2 for the expected number of units allocated). The error bars show confidence intervals for the performance of the deterministic allocations in the support of the leximin allocation. Specifically, we remove the best (resp. the worst) deterministic allocations with an aggregate probability of at most 0.1 from the support, and then measure the best (resp. the worst) performance of any deterministic allocation in the support. A final remark is that the size of the support of the leximin allocation is less than 8 on average in our simulations. A randomization over at most 8 deterministic allocations can easily be stored and implemented in practice, which further supports the practicability of the leximin mechanism.

3.7 Related Work

The problem of fairly dividing a set of indivisible goods has been studied extensively. As an early, seminal example, Hylland and Zeckhauser [113] propose a compelling pseudo-market mechanism to compute a lottery over deterministic assignments, given cardinal preferences. Their mechanism satisfies proportionality, envy-freeness, and ex-ante efficiency, but fails to provide strategyproofness. A more serious objection to their mechanism is that they elicit cardinal utilities from agents — a difficult task in practice. A market approach also drives the work of Budish [44] on approximate competitive equilibrium from equal incomes. His approximation guarantees are practical as long as the supply of each good is relatively large, which is not the case in the classroom allocation setting (where the number of available classrooms in a facility is typically small).

Bogomolnaia and Moulin [29] study random assignment under ordinal preferences. They introduce the *probabilistic serial* (PS) mechanism, which satisfies ex-ante efficiency as well as ordinal fairness. Informally, the probabilistic serial mechanism allows agents to "eat" (at identical speeds) their shares of different goods one by one in the order in which they rank the goods. However, similarly to the pseudo-market mechanism of Hylland and Zeckhauser [113], the probabilistic serial mechanism pertains to the basic setting of assigning n indivisible goods to n agents.

Budish et al. [45] propose a general framework, which, by generalizing the classic Birkhoff von-Neumann theorem [25, 197], extends both mechanisms to handle real-world combinatorial domains, e.g., with group quotas, endogenous capacities, multi-unit non-additive demands, scheduling constraints, etc. Their extension of the probabilistic serial mechanism would be a potential starting point in our setting, if we wished to elicit ordinal preferences from the agents. However, note that in our setting a charter school demanding d classrooms must either receive all d classrooms at a single facility or no classrooms at all — this restriction is incompatible with the framework of Budish et al. [45]. There are other extensions of the probabilistic serial mechanism with multi-unit demands [10, 55, 123, 183], but all of them leverage the standard Birkhoff von-Neumann theorem to allocate *at most* d goods to an agent, and cannot ensure that the agent receives exactly d goods (or no goods at all). We consider it an interesting open problem to extend the probabilistic serial mechanism to the classroom allocation setting with ordinal preferences.

3.8 Epilogue and Discussion of Practical Aspects

In January 2015, PUSD asked charter schools to formally report dichotomous preferences, in addition to the free-text preferences submitted through the usual request form. The plan was to evaluate our approach by comparing its output on the collected explicit dichotomous preferences against human-generated allocations based on the free-text preferences. Despite the promising outlook, sadly, in April 2015 the collaboration was terminated by PUSD, for reasons unknown to us. Nonetheless, we were informed that this initiative helped PUSD build a good rapport with local charter schools.

Meanwhile, Mr. Mopatis put us in touch with representatives of the Los Angeles Unified School District (LAUSD) — the largest school district in California, which includes 274 charter schools and over 900 public schools. The possibility of applying our mechanism at such a large scale is particularly exciting, because the advantages over human-generated allocations are likely to be quite stark. Having faced large lawsuits in the recent past, LAUSD expressed enthusiasm about exploring our approach, and as of June 2016, our approach is currently under consideration by LAUSD for deployment.

More generally, the simplicity of the leximin mechanism, and the intuitiveness of the properties of proportionality, envy-freeness, Pareto optimality, and strategyproofness, have made the approach more likely to be adopted. On the other hand, the use of randomization, though absolutely necessary in order to guarantee fairness in allocating indivisible goods such as classrooms, has been a somewhat harder sell. Ironically, this seems to be the result of presenting the mechanism as a "lottery", which makes it easier to comprehend on the one hand, but on the other hand raises negative connotations and legal objections — even though many charter schools use a (straightforward) lottery system to admit students. In terms of lessons learned, it actually seems better to use more technical terms in this context.

In conclusion, redesigning the way California's school districts allocate classrooms to charter schools is a major project with clear societal impact. This chapter presents a detailed technical approach, but deployment of this approach is still in its infancy; we hope to continue working with school districts for years to come.

Chapter 4

Fair Division of Computational Resources

4.1 Introduction

In the previous chapter, motivated by a real-world setting in which limited resources (unused classrooms) needed to be split fairly, we studied the leximin mechanism and provided a comprehensive analysis. Another prominent application of the leximin mechanism is in computing systems, where resource allocation is a fundamental issue because such systems are naturally constrained in terms of CPU time, memory, communication links, and other resources. Often, these resources must be allocated to multiple *agents* with different requirements. Such situations arise, e.g., in operating systems (where the agents can be jobs) or in cloud computing and data centers (where the agents can be users, companies, or software programs representing them). To take one example, *federated clouds* [185] involve multiple agents that contribute resources; the redistribution of these resources gives rise to delicate issues, including fairness as well as incentives for participation and revelation of private information, which must be carefully considered.

Despite the growing need for resource allocation policies that can address these requirements, state-of-the-art systems employ simple abstractions that fall short. For example, as pointed out by Ghodsi et al. [99], Hadoop and Dryad—two of the most widely-used cluster computing frameworks—employ a *single resource* abstraction for resource allocation. Specifically, these frameworks partition the resources into bundles—known as *slots*—that contain fixed amounts of different resources. The slots are then treated as the system's single resource type, and at this point the allocation can be handled using standard techniques that were developed by the systems community. However, in a realistic environment where agents have heterogeneous demands, the single resource abstraction inevitably leads to significant inefficiencies.

Ghodsi et al. [99] suggest a compelling alternative. Their key insight is that even though agents may have heterogeneous demands for resources, their demands can be plausibly assumed to be highly structured, in maintaining a fixed proportion between resource types. For example, if an agent wishes to execute multiple instances of a job that requires 2 CPUs and 1 GB RAM, its demand for these two resources has a fixed ratio of 2. Given 5 CPUs and 1.8 GB RAM, the agent can run only 1.8 instances of its task (note that Ghodsi et al. allow divisible tasks) despite the additional CPU, hence the agent would be indifferent between this allocation and receiving only 3.6 CPUs and 1.8 GB RAM. Preferences over resource bundles that exhibit this proportional structure are known as *Leontief preferences* in the economics literature. There are some positive results on resource allocation under Leontief preferences [156], but more often than not Leontief preferences are drawn upon for negative examples.

In this model, Ghodsi et al. [99] study the leximin mechanism, albeit under a different name: the *dominant resource fairness (DRF)* mechanism. To be consistent with the literature, we also refer to the mechanism as DRF in this chapter. Ghodsi et al. show that DRF satisfies four key desiderata: sharing incentives (SI),[1] envy-freeness (EF), Pareto optimality (PO), and strategyproofness (SP). One can easily check that Leontief preferences satisfy the four conditions from Section 3.3.2. Hence, this result by Ghodsi et al. is in fact implied by our more general analysis from Chapter 3.

Despite the significant step forward made by Ghodsi et al. [99], there are still many key issues that need to be addressed. Is it possible to rigorously extend the DRF paradigm to more expressive settings where agents are weighted? How can we formally tackle settings where agents' demands are indivisible? How does DRF compare to alternative mechanisms when social welfare is a concern? We provide answers to these questions in this chapter.

4.2 The Model

We begin with an intuitive exposition based on an example from Ghodsi et al. [99], and provide a different perspective on this example. Subsequently, we formulate a simple mathematical model, and more rigorously introduce our notations and assumptions.

Intuition and an alternative interpretation of DRF: Consider a system with 9 CPUs, 18 GB RAM, and two agents. Agent 1 wishes to execute a (divisible) task with the demand vector $\langle 1 \text{ CPU}, 4 \text{ GB}\rangle$, and agent 2 has a (divisible) task that requires $\langle 3 \text{ CPU}, 1 \text{ GB}\rangle$. Note that each instance of the task of agent 1 demands $1/9$ of the total CPU and $2/9$ of the total RAM; the task of agent 2 requires $1/3$ of the total CPU and $1/18$ of the total RAM.

The *dominant resource fairness (DRF)* mechanism [99] works as follows. The *dominant resource* of an agent is the resource for which the agent's task requires the largest fraction of total availability. In the example, the dominant resource of agent 1 is RAM, and the dominant resource of agent 2 is CPU. The DRF mechanism seeks to maximize the number of allocated tasks, under the constraint that the fractions of dominant resource that are allocated—called *dominant shares*—are equalized.

[1] This is simply a different name for proportionality. Again, to be consistent with the literature, we also use the term sharing incentives in this chapter.

Returning to the running example, let y and z be the (possibly fractional) quantities of tasks allocated by DRF to agents 1 and 2, respectively; then overall $y + 3z$ CPUs and $4y + z$ GB are allocated, and these quantities are constrained by the availability of resources. The dominant shares are $2y/9$ for agent 1 and $z/3$ for agent 2. The values of y and z can be computed as follows:

$$\begin{array}{ll} \max\ (y, z) & \\ \text{subject to} & y + 3z \leqslant 9 \\ & 4y + z \leqslant 18 \\ & \frac{2y}{9} = \frac{z}{3} \end{array}$$

Due to the equality $2y/9 = z/3$ it is sufficient to maximize either y or z; we maximize the pair (y, z) for consistency with [99]. The solution is $y = 3$ and $z = 2$, i.e., agent 1 is allocated 3 CPUs and 12 GB RAM, and agent 2 is allocated 6 CPUs and 2 GB RAM.

We next introduce a novel, somewhat different way of thinking about DRF, which greatly simplifies the analysis of its properties. Let D_{ir} be the ratio between the demand of agent i for resource r, and the availability of that resource. In our example, when CPU is resource 1 and RAM is resource 2, $D_{11} = 1/9$, $D_{12} = 2/9$, $D_{21} = 1/3$, and $D_{22} = 1/18$. For all agents i and resources r, denote $d_{ir} = D_{ir}/(\max_{r'} D_{ir'})$; we refer to these demands as *normalized demands*. In the example, $d_{11} = 1/2$, $d_{12} = 1$, $d_{21} = 1$, $d_{22} = 1/6$.

We propose a linear program whose solution x is the dominant share of each agent. Observing the fixed proportion between an agent's demands for different resources, agent i is allocated an $(x \cdot d_{ir})$-fraction of resource r.

$$\begin{array}{ll} \max\ x & \\ \text{subject to} & \sum_i x \cdot d_{ir} \leqslant 1, \quad \forall r \end{array}$$

Allocations are bounded from above by 1 because the program allocates fractions of the total availability of resources. Clearly this linear program can be rewritten as:

$$x = \frac{1}{\max_r \sum_i d_{ir}}. \tag{4.1}$$

In the example, $x = 1/(1/2 + 1) = 2/3$. That is, the allocation to agent 1 is 1/3 of the total CPU, and 2/3 of the total RAM, which is equivalent to 3 CPUs and 12 GB RAM, as before. Similarly, agent 2 is allocated 2/3 of the total CPU and 1/9 of the total RAM, which is equivalent to 6 CPUs and 2 GB RAM, as before.

Rigorous model: Denote the set of agents by $N = \{1, \ldots, n\}$, and the set of resources by R, $|R| = m$. As above, we denote the normalized demand vector of agent $i \in N$ by $\boldsymbol{d}_i = \langle d_{i1}, \ldots, d_{im} \rangle$, where $0 \leqslant d_{ir} \leqslant 1$ for all $r \in R$. An allocation $\boldsymbol{A}$ allocates a fraction A_{ir} of resource r to agent i, subject to the feasibility condition $\sum_{i \in N} A_{ir} \leqslant 1$ for all $r \in R$. A resource allocation *mechanism* is a function that receives normalized demand vectors as input, and outputs an allocation.

Throughout the chapter we assume that resources (e.g., CPU, RAM) are divisible. Up to (but excluding) Section 4.5, our model for preferences coincides with the domain of *Leontief preferences*. Let the utility of an agent for its allocation vector $\boldsymbol{A}_i$ be

$$u_i(\boldsymbol{A}_i) = \max\{y \in \mathbb{R}_+ : \ \forall r \in R,\ A_{ir} \geqslant y \cdot d_{ir}\}.$$

In words, the utility of an agent is the fraction of its dominant resource that it can actually use, given its proportional demands and its allocation of the various resources. Note that *this definition makes the implicit assumption that agents' tasks are divisible*; we relax this assumption in Section 4.5. Unless explicitly mentioned otherwise, we *do not* rely on an interpersonal comparison of utilities. Put another way, with the exception of Section 4.4, an agent's utility function simply induces ordinal preferences over allocations, and its exact value is irrelevant.

An allocation $\boldsymbol{A}$ is called *non-wasteful* if for every agent $i \in N$ there exists $y \in \mathbb{R}_+$ such that for all $r \in R$, $A_{ir} = y \cdot d_{ir}$. Note that if $\boldsymbol{A}$ is a non-wasteful allocation then for all $i \in N$,

$$u_i(\boldsymbol{A}'_i) > u_i(\boldsymbol{A}_i) \Rightarrow \forall r \in R \; s.t. \; d_{ir} > 0, \; A'_{ir} > A_{ir}. \tag{4.2}$$

4.3 Extensions: Weights, Zero Demands, and Group Strategyproofness

In this section we depart from the framework of Ghodsi et al. [99] in three ways. First, we allow agents to be weighted, based on their contribution to the system. Second, we explicitly model the case where agents may have zero demands and therefore DRF needs to allocate in multiple rounds (these first two issues were informally considered by Ghodsi et al., as we discussed in Section 4.1). Third, we study stronger game theoretic properties such as group strategyproofness. We show that DRF can be modified to address these realistic extensions, although the analysis of its game theoretic properties becomes more intricate.

For (contribution) weights, we assume that each agent $i \in N$ has a publicly known weight w_{ir} for resource $r \in R$, which reflects the agent's endowment of that resource (think of this as the amount of the resource contributed by the agent to the resource pool or the equivalent monetary contribution made by the agent for the resource). We assume without loss of generality that for all resources $r \in R$, $\sum_{i \in N} w_{ir} = 1$. PO and SP are defined identically, but the weighted setting does require modifications to the notions of SI and EF, which are redefined as follows. SI now means that an agent $i \in N$ receives as much value as it would get from the allocation that assigns it a w_{ir}-fraction of each resource $r \in R$, i.e., $u_i(\boldsymbol{A}_i) \geqslant u_i(\langle w_{i1}, \ldots, w_{im} \rangle)$. EF requires that agent i does not envy agent j when the allocation of j is scaled by w_{ir}/w_{jr} on each resource r. Formally,

$$u_i(\boldsymbol{A}_i) \geqslant u_i(\langle (w_{i1}/w_{j1}) \cdot A_{j1}, \ldots, (w_{im}/w_{jm}) \cdot A_{jm} \rangle).$$

An alternative definition for the weighted version of EF might simply require that in the special case where agents have a uniform weight $w_i = w_{ir}$ for all $r \in R$, agent i does not prefer the allocation of agent j scaled by w_i/w_j. However, note that the above definition is stronger, and, as we shall demonstrate, the stronger version is feasible. We also note that setting $w_{ir} = 1/n$ for all agents $i \in N$ and resources $r \in R$ recovers the *unweighted case*.

Our second contribution in this section is explicitly modeling the case where agents may not demand every resource. At first glance the assumption of positive demands, which implies that DRF can be simulated via a single allocation according to Equation (4.1), seems very mild. After all, it may seem that the allocation when an agent has zero demand for a resource would be very similar to the allocation when that agent has almost zero demand for the same resource.

To see why this is not the case, let the normalized demand vector of agent 1 be $\langle 1, \varepsilon \rangle$ (i.e., $d_{11} = 1$ and $d_{12} = \varepsilon$), and let agents 2 and 3 have the demand vector $\langle \varepsilon, 1 \rangle$. DRF allocates roughly $1/2$ of resource 1 to agent 1, and roughly $1/2$ of resource 2 to each of agents 2 and 3. If the ε demands are replaced by zero, the first-round allocation (using Equation (4.1)) would be similar. However, then a second round would take place; participation is restricted to agents that only demand non-saturated resources. In this second round agent 1 would receive the unallocated half of resource 1, hence agent 1 would ultimately be allocated all, instead of just half, of resource 1. Crucially, note that agent 1 ends up receiving all of its dominant resource while agents 2 and 3 receive $1/2$ of their dominant resources. Thus, agents might receive unequal shares of their dominant resources at the end of such a multi-round DRF allocation in the presence of zero demands, unlike the single-round DRF allocation in the case of non-zero demands.

In dealing with both weights and zero demands, we extend DRF as follows. Let r_i^* be the weighted dominant resource of agent $i \in N$, defined by $r_i^* \in \arg\min_{r \in R:\, d_{ir} > 0} (w_{ir}/d_{ir})$. Informally, an agent's demand is now scaled by its weight on each resource. Next, we let $\rho_i = w_{ir_i^*}/d_{ir_i^*}$ be the ratio of weight to demand on the weighted dominant resource of agent $i \in N$; for now we assume, for ease of exposition, that $d_{ir} > 0$ implies $w_{ir} > 0$ (we explain how to drop this assumption later). The extended DRF mechanism proceeds in rounds; let s_{rt} be the surplus fraction of resource $r \in R$ left unallocated at the beginning of round t (so $s_{r1} = 1$ for all $r \in R$). The fraction of resource r allocated to agent i in round t is denoted by A_{irt}. We remark that agent weights and demands do not change during the execution of the mechanism, but an agent's dominant resource might change depending on the surpluses of various resources available in each round.

We present our mechanism EXTENDEDDRF as Algorithm 4. It deviates from the unweighted mechanism in a couple of ways. Instead of using $x = 1/(\max_r \sum_i d_{ir})$ as in Equation (4.1), it defines x_t in round t using the ρ_i's of the agents. Thus, each agent i is now allocated an $(x_t \cdot \rho_i \cdot d_{ir})$-fraction of resource r in round t, and hence an $(x_t \cdot w_{ir_i^*})$-fraction of its dominant resource, instead of being allocated an x-fraction of its dominant resource and other resources in proportion, as in the unweighted case.

Our main result of this section is that EXTENDEDDRF satisfies the four desirable properties introduced above. In fact, the mechanism satisfies an even stronger game-theoretic property known as *group strategyproofness (GSP)*: whenever a coalition of agents misreports demands, there is a member of the coalition that does not strictly gain. An interested reader can check that EXTENDEDDRF is the deterministic and weighted version of the leximin mechanism in the Leontief preference domain. Chapter 3 provides a general analysis for the randomized and unweighted version of the leximin

ALGORITHM 4: EXTENDEDDRF

Data: Demands $\boldsymbol{d}$, weights $\boldsymbol{w}$
Result: An allocation $\boldsymbol{A}$
$t \leftarrow 1$;
$\forall r, s_{r1} \leftarrow 1$;
$S_1 \leftarrow N$;
while $S_t \neq \emptyset$ **do**
  $x_t \leftarrow \min_{r \in R}\left(\frac{s_{rt}}{\sum_{i \in S_t} \rho_i \cdot d_{ir}}\right)$;
  $\forall i \in S_t, r \in R, A_{irt} \leftarrow x_t \cdot \rho_i \cdot d_{ir}$;
  $\forall i \in N \setminus S_t, r \in R, A_{irt} \leftarrow 0$;
  $\forall r \in R, s_{r,t+1} \leftarrow s_{rt} - \sum_{i \in S_t} A_{irt}$;
  $t \leftarrow t + 1$;
  $S_t \leftarrow \{i \in S_{t-1} : \forall r \in R,\ d_{ir} > 0 \Rightarrow s_{rt} > 0\}$; /* S_t demand unsaturated resources */
end
$\forall i \in N, r \in R, A_{ir} \leftarrow \sum_{k=1}^{t-1} A_{irk}$

mechanism. First, note that divisibility of resources makes the Leontief preferences domain convex. Hence, the deterministic and the randomized versions of the leximin mechanism provide identical utilities to the agents. Next, we remark that the analysis in Chapter 3 can easily be extended to the weighted version of the leximin mechanism. Nonetheless, we provide the original proof for EXTENDEDDRF in the Leontief preferences domain published in the conference version of this work [164] for the sake of completeness.

Theorem 4.1. EXTENDEDDRF *is PO, SI, EF, and GSP.*

It is trivial that the mechanism is PO, because resources are allocated in the correct proportions, i.e., it is non-wasteful, and moreover it allocates resources as long as players can derive value from them. Establishing SI is also a simple matter. Indeed, for every $r \in R$,

$$\sum_{i \in N} \rho_i \cdot d_{ir} \leqslant \sum_{i \in N:\, d_{ir} > 0} \left(\frac{w_{ir}}{d_{ir}} \cdot d_{ir}\right) = \sum_{i \in N:\, d_{ir} > 0} w_{ir} \leqslant 1$$

and therefore $x_1 \geqslant 1$. Thus, each agent i receives $x_1 \cdot \rho_i \cdot d_{ir_i^*} = x_1 \cdot w_{ir_i^*} \geqslant w_{ir_i^*}$ of resource r_i^* already in round 1. We need to show that each agent $i \in N$ values its allocation as much as the allocation $\langle w_{i1}, \ldots, w_{im} \rangle$, and this now follows from the fact that the mechanism is non-wasteful and from Equation (4.2).

For EF, let $i, j \in N$; we argue that agent i does not envy j. Indeed, let t_i and t_j, respectively, be the last rounds in which these agents were allocated resources. If $t_j > t_i$, agent j does not demand some resource that i does and hence is not allocated a share of that resource. We can therefore assume that $t_j \leqslant t_i$. If $d_{jr_i^*} > 0$ then

$$A_{jr_i^*} = \left(\sum_{t=1}^{t_j} x_t\right) \cdot \rho_j \cdot d_{jr_i^*} \leqslant \left(\sum_{t=1}^{t_j} x_t\right) \cdot \frac{w_{jr_i^*}}{d_{jr_i^*}} \cdot d_{jr_i^*} = \left(\sum_{t=1}^{t_j} x_t\right) \cdot w_{jr_i^*},$$

and otherwise $A_{jr_i^*} = 0$. Moreover,

$$A_{ir_i^*} = \left(\sum_{t=1}^{t_i} x_t\right) \cdot \rho_i \cdot d_{ir_i^*} = \left(\sum_{t=1}^{t_i} x_t\right) \cdot w_{ir_i^*} \geqslant \left(\sum_{t=1}^{t_j} x_t\right) \cdot w_{ir_i^*}.$$

Scaling $A_{jr_i^*}$ by $(w_{ir_i^*}/w_{jr_i^*})$, we get at most $A_{ir_i^*}$; EF is then implied by Equation (4.2).

The challenge is to show that EXTENDEDDRF is GSP. To gain some insight, let us concentrate first on a very special case: strictly positive demands (one round), no weights ($\rho_i = 1$ for all $i \in N$), and SP rather than GSP. For this, assume that agent $i \in N$ reports the demand vector $\boldsymbol{d}'_i$ instead of $\boldsymbol{d}_i$; this leads to the solution x' to Equation (4.1) instead of x, which induces the allocation $A'_{ir} = x' \cdot d'_{ir}$. If $x' \leqslant x$ then agent i receives $x' \cdot d'_{ir_i} \leqslant x \cdot 1 = x \cdot d_{ir_i}$ of its dominant resource r_i, so its utility cannot increase by Equation (4.2). If $x' > x$, consider a resource r that was saturated when reporting $\boldsymbol{d}_i$. It holds that

$$A_{ir} = 1 - \sum_{j \neq i} A_{jr} = 1 - \sum_{j \neq i} x \cdot d_{jr} > 1 - \sum_{j \neq i} x' \cdot d_{jr} = 1 - \sum_{j \neq i} A'_{jr} \geqslant A'_{ir},$$

hence, once again, the utility of agent i cannot increase by Equation (4.2).

The next lemma extends the above argument to the case of possibly zero demands and weighted agents, and strengthens SP to get GSP, thereby completing the proof of Theorem 4.1.

Lemma 4.1. EXTENDEDDRF *is GSP.*

Proof. Assume, without loss of generality, that the set of manipulating agents is $M \subseteq N$ and they report untruthful normalized demand vectors $\boldsymbol{d}'_M = \langle \boldsymbol{d}'_i \rangle_{i \in M}$ (which induce ratios ρ'_i); let $\boldsymbol{d}'$ be the collection of normalized demand vectors where $\boldsymbol{d}'_i = \boldsymbol{d}_i$ for all $i \in N \setminus M$.

Let t^* be the first round when a demanded resource of some $i \in M$ becomes saturated, under truthful or untruthful demands (i.e., the minimum of the two rounds). Let *isSaturated* be a Boolean variable that is true if and only if a demanded resource of a manipulator becomes saturated at the end of round t^* under the truthful reports, and define *isSaturated'* similarly for the untruthful reports. As before, denote by A_{irt} (resp., A'_{irt}) the share of resource $r \in R$ allocated to agent $i \in N$ in round t, and let $A_{ir} = \sum_t A_{irt}$ (resp., A'_{ir}) be the total fraction of resource $r \in R$ allocated to agent $i \in N$ under $\boldsymbol{d}$ (resp., $\boldsymbol{d}'$). In addition, S_t (resp., S'_t) is the set of agents that demand only unsaturated resources in round t under $\boldsymbol{d}$ (resp., $\boldsymbol{d}'$).

Claim 4.1. *1. For all $t \leqslant t^*$ and $r \in R$ such that $d_{ir} = d'_{ir} = 0$ for all $i \in M$, $s_{rt} = s'_{rt}$.*
2. For all $t \leqslant t^*$, $S_t = S'_t$.
3. For all $t < t^*$, $x_t = x'_t$.

Proof. We prove the claim by induction on t. The base of the induction is trivial, as $s_{r1} = s'_{r1} = 1$ for all $r \in R$, and $S_1 = S'_1 = N$. We can also let $x_0 = x'_0$ as a formality.

For the induction step, consider round $t < t^*$. We assume statements (1) and (2) are true for all $t' \leqslant t$, and that statement (3) is true for all $t' < t$. We prove statement (3) for round t and then we prove statements (1) and (2) for round $t + 1$.

In round t, no agent $i \in M$ has saturated resources under $\boldsymbol{d}$ or $\boldsymbol{d}'$, hence there are resources $r, r' \in R$ such that for all $i \in M$, $d_{ir} = d'_{ir} = 0$ and $d_{ir'} = d'_{ir'} = 0$, which become saturated at round t under the two demand vectors, i.e.,

$$x_t = \frac{s_{rt}}{\sum_{i \in S_t} \rho_i \cdot d_{ir}},$$

and

$$x'_t = \frac{s'_{r't}}{\sum_{i \in S'_t} \rho'_i \cdot d'_{ir'}}.$$

By the induction assumption, $s_{rt} = s'_{rt}$, $s_{r't} = s'_{r't}$, and $S_t = S'_t$. Moreover, $\sum_{i \in S_t} \rho_i \cdot d_{ir} = \sum_{i \in S'_t} \rho'_i \cdot d'_{ir}$ and $\sum_{i \in S_t} \rho_i \cdot d_{ir'} = \sum_{i \in S'_t} \rho'_i d'_{ir'}$, where last two equalities hold because every $i \in M$ does not demand either r or r' under $\boldsymbol{d}$ or $\boldsymbol{d}'$ (and hence the summations include zero terms for these agents). It follows that

$$x_t = \frac{s_{rt}}{\sum_{i \in S_t} \rho_i \cdot d_{ir}} = \frac{s'_{rt}}{\sum_{i \in S'_t} \rho'_i \cdot d'_{ir}} \geqslant x'_t,$$

and similarly $x'_t \geqslant x_t$. We conclude that $x_t = x'_t$.

To establish statement (1) for round $t+1$, let $r \in R$ such that $d_{ir} = d'_{ir} = 0$ for all $i \in M$. Using the induction assumption we conclude that

$$s_{r,t+1} = s_{rt} - x_t \sum_{i \in S_t} \rho_i d_{ir} = s'_{rt} - x'_t \sum_{i \in S'_t} \rho'_i d'_{ir} = s'_{r,t+1}.$$

Finally, the assertion that $S_{t+1} = S'_{t+1}$ follows from the fact that all resources not demanded by members of the coalition have the same surplus under both $\boldsymbol{d}$ and $\boldsymbol{d}'$, and resources that are demanded by at least one member of the coalition are not saturated under $\boldsymbol{d}$ nor under $\boldsymbol{d}'$. ■

Having established Claim 4.1, we proceed with the lemma's proof by distinguishing between four cases.

Case 1: $x'_{t^*} \geqslant x_{t^*}$ and isSaturated is true. Let r be a resource that is saturated under $\boldsymbol{d}$ in round t^*, and is demanded by some manipulator, i.e., there is $i \in M$ such that $d_{ir} > 0$. Note that such a resource exists by the definition of the variable *isSaturated*. Using Claim 4.1, we have that for every $t < t^*$ and every $i \in N \setminus M$,

$$A_{irt} = x_t \cdot \rho_i \cdot d_{ir} = x'_t \cdot \rho'_i \cdot d'_{ir} = A'_{irt},$$

and similarly $A_{irt^*} \leqslant A'_{irt^*}$ by the assumption that $x_{t^*} \leqslant x'_{t^*}$. It follows that for all $i \in N \setminus M$,

$$A_{ir} = \sum_{t=1}^{t^*} A_{irt} \leqslant \sum_{t=1}^{t^*} A'_{irt} \leqslant A'_{ir}.$$

Consider the set of agents $M' = \{i \in M | d_{ir} > 0\}$. From the definition of r, M' is nonempty. In addition, for all $i \in M \backslash M'$, $A_{ir} = 0 \leqslant A'_{ir}$. We conclude that for any $i \in N \setminus M'$, $A_{ir} \leqslant A'_{ir}$. It follows that

$$\sum_{i \in M'} A'_{ir} \leqslant 1 - \sum_{i \in N \backslash M'} A'_{ir} \leqslant 1 - \sum_{i \in N \backslash M'} A_{ir} = \sum_{i \in M'} A_{ir},$$

where the equality holds because r is saturated under $\boldsymbol{d}$. In other words, the overall allocation of resource r to the set of manipulating agents in M' does not grow larger as a result of manipulation, so there must be $i \in M'$ such that $A'_{ir} \leqslant A_{ir}$. Since $d_{ir} > 0$, this implies by Equation (4.2) that the utility to agent i does not increase, as required.

Case 2: $x'_{t^*} \leqslant x_{t^*}$ and isSaturated′ is true. Consider an agent $i \in M$ that demands a resource that becomes saturated in round t^* under $\boldsymbol{d}'$. Let r_i^* be a weighted dominant resource of agent i. We have that $\rho_i = w_{ir_i^*}/d_{ir_i^*}$ and $\rho'_i \leqslant w_{ir_i^*}/d'_{ir_i^*}$. It follows that $\rho_i \cdot d_{ir_i^*} = w_{ir_i^*}$ and $\rho'_i \cdot d'_{ir_i^*} \leqslant w_{ir_i^*}$. From Claim 4.1, we know that for all $t < t^*$, $x_t = x'_t$, and our assumption is that $x'_{t^*} \leqslant x_{t^*}$. Therefore, for every $t \leqslant t^*$,

$$A'_{ir_i^*t} = x'_t \cdot \rho'_i \cdot d'_{ir_i^*} \leqslant x'_t \cdot w_{ir_i^*} \leqslant x_t \cdot w_{ir_i^*} = x_t \cdot \rho_i \cdot d_{ir_i^*} = A_{ir_i^*t}.$$

Next we notice that under $\boldsymbol{d}'$, agent i is not allocated more resources after round t^*. We conclude that

$$A'_{ir_i^*} = \sum_{t=1}^{t^*} A'_{ir_i^*t} \leqslant \sum_{t=1}^{t^*} A_{ir_i^*t} \leqslant A_{ir_i^*}.$$

Thus, there exists an agent $i \in M$ who does not receive an increased share of weighted dominant resource under $\boldsymbol{d}'$, and hence its utility under $\boldsymbol{d}'$ does not increase.

Case 3: $x'_{t^*} > x_{t^*}$ and isSaturated is false (and hence isSaturated′ is true). We argue that this case is impossible. Let r be a resource that is saturated under $\boldsymbol{d}$ in round t^*. Since *isSaturated* is false, it must hold that for all $i \in M$, $d_{ir} = 0$. Thus for all $t \leqslant t^*$ (using $S_t = S'_t$ by Claim 4.1), $\sum_{i \in S'_t} \rho'_i \cdot d'_{ir} \geqslant \sum_{i \in S_t} \rho_i \cdot d_{ir}$. Claim 4.1 further states that $x_t = x'_t$ for all $t < t^*$, and therefore

$$s_{rt^*} = 1 - \sum_{t=1}^{t^*-1} \left(x_t \cdot \sum_{i \in S_t} \rho_i \cdot d_{ir} \right) \geqslant 1 - \sum_{t=1}^{t^*-1} \left(x'_t \cdot \sum_{i \in S'_t} \rho'_i \cdot d'_{ir} \right) = s'_{rt^*}.$$

We conclude that

$$x_{t^*} = \frac{s_{rt^*}}{\sum_{i \in S_{t^*}} \rho_i \cdot d_{ir}} \geqslant \frac{s'_{rt^*}}{\sum_{i \in S'_{t^*}} \rho'_i \cdot d'_{ir}} \geqslant x'_{t^*},$$

Which contradicts our assumption that $x'_{t^*} > x_{t^*}$.

Case 4: $x'_{t^*} < x_{t^*}$ and isSaturated′ is false (and hence isSaturated is true). This case is symmetric to case 3, by replacing the roles of $\boldsymbol{d}$ and $\boldsymbol{d}'$.

It may not be immediately apparent that the above four cases are exhaustive, but note that it is never the case that both *isSaturated* and *isSaturated′* are false. Therefore, the three possible combinations of values for *isSaturated* and *isSaturated*$'$ are covered by cases 1 and 3 when $x'_{t^*} > x_{t^*}$; by cases 2 and 4 when $x'_{t^*} < x_{t^*}$; and by cases 1 and 2 when $x'_{t^*} = x_{t^*}$. This completes the proof of Lemma 4.1. ■

While GSP is a strong game-theoretic axiom, an even stronger version has been studied in the literature. Under the stronger notion, it cannot be the case that the value of all manipulators is at least as high, and the value of at least one manipulator is strictly higher under the false reported demands. Under the assumption of strictly positive demands, it is easy to verify that EXTENDEDDRF also satisfies the stronger notion (the allocation is made in a single round).

However, it turns out that when there are multiple rounds, one agent can help another by causing a third agent to drop out early without losing value itself. To see this, consider the following setting with 5 agents with $w_{ir} = 1/5$ for each $i \in N$ and $r \in R$, and,

$$\boldsymbol{d_1} = \langle 0,1,0\rangle, \boldsymbol{d_2} = \langle 1,0,0\rangle, \boldsymbol{d_3} = \langle 1,0,1/4\rangle, \boldsymbol{d_4} = \langle 0,1,1\rangle, \boldsymbol{d_5} = \langle 0,1,1\rangle.$$

It can be seen that $x_1 = 1/3$. The surplus at the beginning of the second round is $s_{12} = 1/3$, $s_{22} = 0$, and $s_{32} = 1/4$. Resource 2 is saturated and $S_2 = \{2,3\}$. In the second round, $x_2 = 1/6$. Therefore, the utilities of agents 1 and 2 in this setting are $u_1 = 1/3$ and $u_2 = 1/3 + 1/6 = 1/2$.

Now, assume that agents 1 and 2 collude, and report demand vectors $\boldsymbol{d'_1} = \{0,1,3/4\}$ and $\boldsymbol{d'_2} = \boldsymbol{d_2}$. In this case $x'_1 = 1/3$ as before, but the surplus at the beginning of the second round is $s'_{12} = 1/3$, $s'_{22} = 0$, and $s'_{32} = 0$. Hence, both resources 2 and 3 are saturated and $S_2 = \{2\}$. In round 2 agent 2 receives the remaining surplus of resource 1. Hence, under the false demands, $u'_1 = 1/3$ and $u'_2 = 1/3 + 1/3 = 2/3$; agent 1 is not worse off and agent 2 is better off.

We conclude with two remarks. First, note that in each round, the mechanism exhausts a resource that at least one agent demanded. This gives an immediate upper bound of $\min(|N| = n, |R| = m)$ on the number of rounds. Since each round takes $O(n \cdot m)$ time to execute, it follows that the running time of the mechanism is polynomial in the number of agents and the number of resources.

Second, recall that we have assumed for ease of exposition that $d_{ir} > 0$ implies $w_{ir} > 0$. We now briefly explain how to drop this assumption. Observe that if there exist $i \in N$ and $r \in R$ such that $d_{ir} > 0$ but $w_{ir} = 0$, agent i is not entitled to anything according to EF or SI, i.e., any allocation would satisfy these two properties with respect to i. Hence, we can initially remove such agents, and proceed as before with the remaining agents. We then add a second stage where the remaining resources are allocated to the agents that were initially removed, e.g., via unweighted DRF. Using Theorem 4.1, it is easy to verify that this two-stage mechanism is PO, EF, SI, and GSP.

4.4 Limitations: Strategyproof Mechanisms and Welfare Maximization

In Section 4.3 we have established that the DRF paradigm is robust to perturbations of the model, maintaining its many highly desirable properties. In this section we examine some of the limitations of DRF, and ask whether these limitations can be circumvented.

Throughout this section we assume that an agent's utility is exactly $u_i(\boldsymbol{A}_i) = \max\{y \in \mathbb{R}_+ : \forall r \in R,\ A_{ir} \geqslant y \cdot d_{ir}\}$. In particular, if the mechanism is non-wasteful, an agent's utility is the allocated fraction of its dominant resource. In addition, we assume an interpersonal comparison of utilities; i.e., the utility function is not merely a formalism for comparing allocations ordinally. Therefore, given an allocation $\boldsymbol{A}$, we define its (utilitarian) social welfare as $\sum_{i \in N} u_i(\boldsymbol{A}_i)$. The results of this section, which are negative in nature, hold under this assumption specifically, but similar results hold under perturbations of the utility functions such as linear affine sums of utility. Because the negative results hold even when players are unweighted in terms of their resource contributions, we adopt the model and definitions of Section 4.2.

Our basic observation is that DRF may provide very low social welfare compared to a welfare-maximizing allocation. For example, consider a setting with m resources. For each resource r there is an agent i_r such that $d_{i_r,r} = 1$, and $d_{i_r,r'} = 0$ for all $r' \in R \setminus \{r\}$. Moreover, there is a large number of agents with normalized demands $d_{ir} = 1$ for all $r \in R$. The optimal allocation would give all of resource r to agent i_r, for a social welfare of m. In contrast, under DRF each agent would receive a $(1/(n-m+1))$-fraction of its dominant resource. As n grows larger, the social welfare of $n/(n-m+1)$ approaches 1. The ratio between these two values is arbitrarily close to m. More formally, letting the *approximation ratio* of a mechanism be the worst-case ratio between the optimal solution and the mechanism's solution, we say that for every $\delta > 0$, DRF cannot have an approximation ratio better than $m - \delta$ for the social welfare.

In some sense it is not surprising that DRF provides poor guarantees with respect to social welfare. After all, the DRF paradigm focuses on *egalitarian welfare* by equalizing the dominant shares across the agents, rather than aiming for utilitarian welfare. So, can other mechanisms do better?

Let us first examine mechanisms that satisfy SI. It is immediately apparent that the above example provides a similar lower bound. Indeed, any SI mechanism would have to allocate at least $1/n$ of each resource to each of the $n-m$ agents with an all-ones normalized demand vector. Intuitively, an SI mechanism must allocate almost εm of the various resources in order to obtain ε social welfare. That is, for any $\delta > 0$, *any* SI mechanism cannot have an approximation ratio better than $m - \delta$ for the social welfare. A similar (though slightly more elaborate) argument works for EF, by replacing the demand vectors of the m agents i_r with $d_{i_r,r} = 1$ (as before) and $d_{i_r,r'} = \varepsilon$ for all $r' \in R \setminus \{r\}$ and an arbitrarily small $\varepsilon > 0$.

These observations are disappointing, but not entirely unexpected. SI and EF are properties that place significant constraints on allocations. SP is an altogether different matter. *A priori* the constraints imposed by SP seem less obstructive to welfare maximization than SI or EF, and indeed in some settings SP mechanisms (even without the use of payments) provide optimal, or nearly optimal, social welfare [174]. Nevertheless, the main result of this section is a similar lower bound for SP mechanisms.

Theorem 4.2. *For any $\delta > 0$ there exists a sufficiently large number of agents n such that no SP mechanism can provide an approximation ratio smaller than $m - \delta$ for the social welfare.*

Proof. We first introduce some notation. Given an allocation $\boldsymbol{A}$, define the *index set* of $\boldsymbol{A}$ as $I_{\boldsymbol{A}} = \{i \in N \mid \min_{r \in R} A_{ir} \leqslant 1/\sqrt{n}\}$. In words, $I_{\boldsymbol{A}}$ denotes the set of agents who have at most a $(1/\sqrt{n})$-fraction of some resource allocated to them. We argue that for any allocation $\boldsymbol{A}$,

$$|I_{\boldsymbol{A}}| \geqslant n - \sqrt{n}. \tag{4.3}$$

Indeed, for any $i \in N \backslash I_{\boldsymbol{A}}$ and $r \in R$, $A_{ir} > 1/\sqrt{n}$. Hence, there cannot be more than $\sqrt{n}$ such agents, otherwise their total allocation for every resource would be more than 1.

Let $\mathbf{1} = (1, 1, \ldots, 1)$ be a the all-ones demand vector. Let $\boldsymbol{\varepsilon}_i$ be the demand vector that has 1 in i^{th} position and ε everywhere else. We are interested in settings where the agents in a subset $X \subseteq N$, $|X| = m$, have the demand vectors $\boldsymbol{\varepsilon}_1, \boldsymbol{\varepsilon}_2, \ldots, \boldsymbol{\varepsilon}_m$ and the remaining agents have demand vectors $\mathbf{1}$. We show that if the resulting allocation A is such that for all $i \in X$, $i \in I_{\boldsymbol{A}}$, then the theorem follows.

Claim 4.2. *Suppose that for any sufficiently large $n \in \mathbb{N}$, for any SP mechanism for n agents and m resources, and for any $\varepsilon > 0$, there exists $X \subseteq N$ such that $|X| = m$ and when the agents in X report the demand vectors $\boldsymbol{\varepsilon}_1, \ldots, \boldsymbol{\varepsilon}_m$, and the agents in $N \setminus X$ report the demand vector $\mathbf{1}$, the mechanism returns an allocation $\boldsymbol{A}$ such that $X \subseteq I_{\boldsymbol{A}}$. Then for any $\delta > 0$, no SP mechanism for m resources can provide an approximation ratio smaller than $m - \delta$.*

Proof. Given n, an SP mechanism, and ε, let $X \subseteq N$ as in the claim's statement. For all $i \in X$, $u_i(\boldsymbol{A}_i) \leqslant 1/(\sqrt{n} \cdot \varepsilon)$. Clearly it also holds that

$$\sum_{i \in N \backslash X} u_i(\boldsymbol{A}_i) \leqslant 1.$$

Hence, the social welfare of an SP mechanism on this instance can be at most $m/(\sqrt{n} \cdot \varepsilon) + 1$.

In contrast, maximum social welfare is obtained by dividing the resources among the m agents in X, with everyone getting equal shares of their dominant resources. Under this allocation, every agent $i \in X$ receives a $(1/(1 + (m-1) \cdot \varepsilon))$-fraction of its dominant resource, which is its utility. Hence, the optimal social welfare is $m/(1 + (m-1) \cdot \varepsilon)$.

It follows that the approximation ratio of the SP mechanism cannot be smaller than

$$\frac{m}{1 + (m-1) \cdot \varepsilon} \cdot \frac{1}{\frac{m}{\sqrt{n} \cdot \varepsilon} + 1}. \tag{4.4}$$

To prove the claim it remains to show that for any $\delta > 0$, we can choose n and ε such that the expression in Equation (4.4) is greater than $m - \delta$. Indeed, choose

$$\varepsilon < \frac{\delta}{3 \cdot m \cdot (m-1)}$$

and

$$n > \max\left(\left(\frac{3 \cdot m^2 \cdot (m-1)}{\delta}\right)^2, \left(\frac{3 \cdot m^2}{\varepsilon \cdot \delta}\right)^2\right).$$

Then $(m^2 \cdot (m-1))/\sqrt{n} < \delta/3$, $m \cdot (m-1) \cdot \varepsilon < \delta/3$, and $m^2/(\sqrt{n} \cdot \varepsilon) < \delta/3$. It follows that

$$\delta > \frac{m^2 \cdot (m-1)}{\sqrt{n}} + m \cdot (m-1) \cdot \varepsilon + \frac{m^2}{\sqrt{n} \cdot \varepsilon},$$

and in particular

$$1 + \frac{\delta}{m} > (1 + (m-1) \cdot \varepsilon) \cdot \left(\frac{m}{\sqrt{n} \cdot \varepsilon} + 1\right).$$

Therefore, using the fact that $1/(1 - \frac{\delta}{m}) > 1 + \delta/m$,

$$\frac{m}{m - \delta} > (1 + (m-1) \cdot \varepsilon) \cdot \left(\frac{m}{\sqrt{n} \cdot \varepsilon} + 1\right),$$

and finally

$$m - \delta < \frac{m}{1 + (m-1) \cdot \varepsilon} \cdot \frac{1}{\frac{m}{\sqrt{n} \cdot \varepsilon} + 1},$$

as required. ■ (Claim 4.2)

Our goal is therefore to prove the hypothesis made in the statement of Claim 4.2. In the remainder of the proof ε is arbitrary and fixed, and we vary the number of agents n. Since we are dealing with an arbitrarily large n, we can assume that m divides n for ease of exposition. We start from a setting where the demand vector is $\boldsymbol{d}_i = \mathbf{1}$ for all $i \in N$. We partition the set of n agents into m buckets, $B_1, B_2, \ldots, B_m$, each with n/m agents. For consistency we impose the following restriction: an agent in B_i can only change its demand from $\mathbf{1}$ to $\boldsymbol{\varepsilon}_i$. In future references, we omit the unambiguous initial and final demands when we say that an agent changes its demand. We show not only that a set $X \subseteq N$ of m agents exists as in the statement of Claim 4.2, but one such set exists that consists of one agent from each bucket, such that after these m agents change their demands, they all belong to the index set of the resulting allocation.

Call an m-tuple $(i_1, \ldots, i_m)$ *diverse* if $i_k \in B_k$ for every $k \in \{1, \ldots, m\}$, that is, if it consists of one agent from each bucket. Let $T_{n,m}$ denote the set of all diverse m-tuples. Let $L_{n,m} \subseteq T_{n,m}$ be the set of diverse m-tuples such that when they all change their demands, at least one of them is not in the index set. We want to show that $|T_{n,m} \setminus L_{n,m}| > 0$. We prove the following claim.

Claim 4.3. *For any $n, m \in \mathbb{N}$, $|T_{n,m}| = (n/m)^m$ and $|L_{n,m}| \leqslant m \cdot (n/m)^{m-1} \cdot \sqrt{n}$.*

Proof. Since $T_{n,m}$ is the set of all diverse m-tuples where one agent is selected from each of m buckets with each bucket containing n/m agents, it is easy to see that $|T_{n,m}| = (n/m)^m$. Recall that $L_{n,m}$ denotes the set of all diverse m-tuples such that when they change their demands, at least one of them is not in the index set. For any $k \in \{1, \ldots, m\}$, define $L^k_{n,m} \subseteq L_{n,m}$ to be the set of m-tuples in $L_{n,m}$ such that when they change their demands, the agent from bucket B_k is not in the index set. Clearly $|L_{n,m}| \leqslant \sum_{k=1}^{m} |L^k_{n,m}|$.

Fix any $k \in \{1, \ldots, m\}$. We upper-bound $|L^k_{n,m}|$ as follows. Take any tuple $t = (i_1, \ldots, i_m) \in L^k_{n,m}$. Let $t_{-k} = (i_1, \ldots, i_{k-1}, i_{k+1}, \ldots, i_m)$. If all agents in t change their demands, agent i_k is not in the index set and gets more than a $(1/\sqrt{n})$-fraction of every resource. Consider the case when only agents in t_{-k} change their demands. If agent i_k is in the index set, then it receives at most a $(1/\sqrt{n})$-fraction of some resource. Thus, its utility under the demand vector $\mathbf{1}$ is at most $1/\sqrt{n}$ and it has strict incentive to misreport its demand vector (to $\boldsymbol{\varepsilon}_k$) and be outside the index set. This is impossible since the mechanism is SP. Thus, i_k must not be in the index set when the agents in t_{-k} change their demands.

Now, for any fixed t_{-k}, the number of possible agents i_k that are not in the index set when the agents in t_{-k} change their demands is at most $\sqrt{n}$ (from Equation (4.3)). Thus, for each possible t_{-k}, there are at most $\sqrt{n}$ tuples in $L^k_{n,m}$. Since the number of ways of choosing t_{-k} is at most $(n/m)^{m-1}$, we have that $|L^k_{n,m}| \leqslant (n/m)^{m-1} \cdot \sqrt{n}$. Hence,

$$|L_{n,m}| \leqslant \sum_{k=1}^{m} |L^k_{n,m}| \leqslant m \cdot \left(\frac{n}{m}\right)^{m-1} \cdot \sqrt{n},$$

as required. ■ (Claim 4.3)

Using Claim 4.3, we see that

$$|T_{n,m} \setminus L_{n,m}| \geqslant \left(\frac{n}{m}\right)^{m} - m \cdot \left(\frac{n}{m}\right)^{m-1} \cdot \sqrt{n} = \left(\frac{n}{m}\right)^{m-1} \cdot \sqrt{n} \cdot \left(\frac{\sqrt{n}}{m} - m\right).$$

Thus, for $n > m^4$, we have that $|T_{n,m} \setminus L_{n,m}| > 0$, as required. ■ (Theorem 4.2)

Despite their negative flavor, the results of this section can be seen as vindicating DRF as far as social welfare maximization is concerned. Indeed, asking for just one of the three properties (SI, EF, or SP) seems like a minimal requirement, but already leads to an approximation ratio that is as bad as the one provided by DRF itself (or EXTENDEDDRF, for that matter).

4.5 Indivisible Tasks

So far, our theoretical analysis treated agents' tasks as divisible: if to run a task an agent needs 2 CPUs and 2 GB RAM, but it is allocated 1 CPU and 1 GB RAM, then it can run half a task. This assumption, which coincides with Leontief preferences, is the driving force behind the results of Section 4.3 as well as earlier and related results [97, 99, 131].

However, in practice agents' tasks would usually be indivisible. Indeed, this is the case in the implementation and simulations carried out by Ghodsi et al. [99]. In other words, a more realistic domain model is a setting with divisible goods but indivisible tasks: an agent can only derive utility from tasks that are allocated enough resources to run.[2] Hence, informally, an agent's utility function is a step function that increases with every additional instance of its task that it can run. In this section we will *not* require interpersonal comparison of utilities; i.e., as in Section 4.3, a utility function is used only to induce *ordinal* preferences over allocations. However, for ease of exposition we do assume here that agents are unweighted (or, equivalently, all agents have equal weights).

More formally, each agent $i \in N$ now reports a demand bundle $\boldsymbol{b}_i$, where b_{ir} denotes the fraction of resource $r \in R$ that agent i requires to complete one instance of its task. Given an allocation $\boldsymbol{A}$, the utility function of agent $i \in N$ is given by

$$u_i(\boldsymbol{A}_i) = \max\{t \in \mathbb{N} \cup \{0\} : \forall r \in R,\ A_{ir} \geqslant t \cdot b_{ir}\},$$

that is, an agent's utility increases linearly with the number of *complete* instantiations of its task that it can run.

Our positive results focus, as before, on non-wasteful mechanisms that do not allocate resources to agents that cannot use them (our negative results do not make this assumption). For a non-wasteful mechanism we can simply denote by $x_i^{\boldsymbol{A}} = x_i$ the number of bundles allocated to agent $i \in N$ under the allocation $\boldsymbol{A}$, in which case the utility of agent i is $u_i(\boldsymbol{A}) = x_i^{\boldsymbol{A}}$. As a last piece of new notation, we define inequalities between two vectors as follows. We write $\boldsymbol{v} \geqslant \boldsymbol{w}$ if the inequality is satisfied pointwise. Furthermore, $\boldsymbol{v} > \boldsymbol{w}$ if $\boldsymbol{v} \geqslant \boldsymbol{w}$ and there is at least one coordinate k such that $v_k > w_k$.

4.5.1 Impossibility results

Observe that the definitions of PO, SI, EF, and SP are identical to the ones given in Section 4.2. Interestingly, DRF still satisfies SI, EF, and SP under the new "truncated" utilities. However, DRF is no longer PO. For example, say that there are two agents and one resource, and $b_{11} = 1/10$, $b_{21} = 2/5$. DRF would allocate $1/2$ of the resource to each agent, but allocating $3/5$ of the resource to agent 1 and $2/5$ of the resource to agent 2 would be better for agent 1 and equally good for agent 2. An implementation of DRF that allocates tasks sequentially to agents with currently minimum dominant share [99] also suffers from several theoretical flaws when analyzed carefully, e.g., it does not satisfy SI. Note that it is trivial to satisfy SI, EF, and SP without PO, by giving each agent exactly its proportional share of each resource.

The next few results imply that we cannot hope to tweak DRF to achieve all four desirable properties (PO, SI, EF, and SP) under indivisibilities.

Theorem 4.3. *Under indivisibilities there is no mechanism that satisfies PO, SI, and SP.*

Proof. Consider a setting with two agents and a single resource. Both agents have bundles $b_{11} = b_{21} = 1/2 + \varepsilon$, for a small $\varepsilon > 0$. Assume for contradiction that there exists a mechanism that satisfies PO, SI, and SP. By PO, the mechanism allocates one bundle to exactly one of the two agents, without loss of generality agent 1. Now suppose that

[2]The combination of divisible goods and indivisible tasks may lead to an additional, practical issue, in regard to resource fragmentation, which we discuss in Section 4.7.

agent 2 reports $b'_{21} = 1/2$; then any SI allocation must allocate at least $1/2$ of the resource to agent 2, hence the only allocation that is both PO and SI allocates all of the resource to agent 2. This is a contradiction to SP. ■

An even more fundamental issue is that PO and EF are trivially incompatible under indivisibilities. For example, when two agents have a task that requires all available resources, the only (two) Pareto optimal allocations give everything to one of the agents, but these allocations are not EF.

Fortunately there is precedent to studying relaxed notions of EF in related settings. A recent example is given by the work of Budish [44] on the *combinatorial assignment* problem. Budish deals with a related resource allocation setting that models allocations of seats in university courses. Each resource (seats in course) has an integer availability, and agents have preferences over bundles of resources. Budish proposes an approximate version of the notion of *competitive equilibrium from equal incomes* [195].[3] An interesting notion of approximate fairness satisfied by his solution is *envy bounded by a single good*: there exists some good in the allocation of agent $j \in N$ such that $i \in N$ does not envy j when that good is removed.[4] This notion has a natural equivalent in our setting, which we formalize below.

Definition 4.1. A mechanism is *envy-free up to one bundle* (EF1) if for every vector of reported bundles $\boldsymbol{b}$ it outputs an allocation $\boldsymbol{A}$ such that for all $i, j \in N$, $u_i(\boldsymbol{A}_i) \geqslant u_i(\boldsymbol{A}_j - \boldsymbol{b}_i)$.

Subtraction between vectors is pointwise. In words, a mechanism is EF1 if each agent $i \in N$ does not envy agent $j \in N$ if one instance of the task of agent i is removed from the allocation of j. For non-wasteful mechanisms, this is equivalent to $x_j\boldsymbol{b}_j \not\geq (x_i + 2)\boldsymbol{b}_i$ for all $i, j \in N$. Unlike the strict notion of EF, we shall see that the relaxed notion is compatible with PO. However, the following straightforward result implies that the latter notion is incompatible with PO and SP.

Theorem 4.4. *Under indivisibilities there is no mechanism that satisfies PO, EF1, and SP.*

Proof. As above, consider a setting with two agents and one resource. Both agents have bundles $b_{11} = b_{21} = 1/3$. The only feasible allocations that are EF1 and PO allocate $2/3$ of the resource to one agent and $1/3$ to the other. Assume without loss of generality that agent 1 receives two bundles. If agent 2 reports $b'_{21} = 1/6$ then the only PO and EF1 allocation gives $2/3$ of the resource to agent 2 and $1/3$ to agent 1, violating SP. ■

We remark that Theorems 4.3 and 4.4 can easily be extended to multiple resources by having every agent demand all the resources equally. To summarize, the combination SP+SI+EF1 is trivial, but if we add PO and insist on SP we immediately run into impossibilities with either SI or EF1. Unfortunately, SP seems to preclude reasonable

[3] Specifically, he allows the market to clear with an error, and also slightly perturbs the initially equal budgets of agents.

[4] Related notions appear in earlier work, for example the approximate notion of equal treatment of equals of Moulin and Stong [152] and an approximate notion of envy-freeness studied by Lipton et al. [134].

mechanisms when tasks are indivisible. We therefore focus on the other three properties in the context of task indivisibilities, namely PO, SI, and EF1, which, as we shall see, induce a rich axiomatic framework and lead to the design of practical mechanisms.

4.5.2 Sequential Minmax

To design a mechanism that is PO, SI, and EF1 under indivisible tasks, we first introduce a few notations. Given a non-wasteful allocation $\boldsymbol{A}$, let $\text{MaxDom}(\boldsymbol{A}) = \max_{i \in N} \max_{r \in R} A_{ir}$ be the maximum dominant share allocated to an agent. We also let $\boldsymbol{A} \uparrow i$ be the allocation obtained by starting from the allocation $\boldsymbol{A}$ and giving agent $i \in N$ another bundle $\boldsymbol{b}_i$.

We are now ready to present our mechanism, SEQUENTIALMINMAX, which is formally given as Algorithm 5. The mechanism sequentially allocates one bundle at each step to an agent that minimizes the maximum dominant share after allocation.

ALGORITHM 5: SEQUENTIALMINMAX

Data: Bundles $\boldsymbol{b}$
Result: An allocation $\boldsymbol{A}$
$k \leftarrow 1$; $\boldsymbol{A}^0 \leftarrow \mathbf{0}$; $T_1 \leftarrow N$;
while $T_k \neq \emptyset$ **do**
  $M_k \leftarrow \{i \in T_k \mid \forall j \in T_k, \text{MaxDom}(\boldsymbol{A}^{k-1} \uparrow i) \leqslant \text{MaxDom}(\boldsymbol{A}^{k-1} \uparrow j)\}$;
  $i \leftarrow$ any agent in M_k;
  $\boldsymbol{A}^k \leftarrow \boldsymbol{A}^{k-1} \uparrow i$;
  $T_{k+1} \leftarrow \{i \in T_k \mid \boldsymbol{A}^k \uparrow i \text{ is feasible}\}$;
  $k \leftarrow k+1$;
end
return $\boldsymbol{A}^{k-1}$;

Our main result of this section is the following theorem.

Theorem 4.5. SEQUENTIALMINMAX *satisfies PO, SI, and EF1.*

Note that the more intuitive alternative of maximizing the minimum dominant share does not achieve the same properties, nor do variations that consider dominant shares before rather than after allocation.

To establish PO we need to prove that we are allocating a bundle to some agent as long as there exists an agent to which we can allocate; this follows trivially from the mechanism itself since $T_k \neq \emptyset$ implies $M_k \neq \emptyset$. We establish SI and EF1 in the following two lemmas.

Lemma 4.2. SEQUENTIALMINMAX *satisfies SI.*

Proof. Let $\boldsymbol{A}^{\text{SI}}$ be the minimal SI allocation that is obtained by giving each agent $1/n$ of each resource, and then taking back resources that agents cannot use. We show that $\boldsymbol{A}^{\text{SI}} = \boldsymbol{A}^k$ for some k during the execution of SEQUENTIALMINMAX. This is sufficient because subsequent allocations can only increase the shares of resources that agents obtain.

We prove, by induction, that $\boldsymbol{A}^k \leqslant \boldsymbol{A}^{\mathrm{SI}}$ (where allocations are treated as vectors and the inequality is pointwise) until it becomes equal to $\boldsymbol{A}^{\mathrm{SI}}$ for the first time. Note that initially $\boldsymbol{A}^0 \leqslant \boldsymbol{A}^{\mathrm{SI}}$. We will show that until $\boldsymbol{A}^k = \boldsymbol{A}^{\mathrm{SI}}$, the mechanism would only allocate a bundle to an agent that is also allocated in $\boldsymbol{A}^{\mathrm{SI}}$, and therefore after finitely many iterations $\boldsymbol{A}^k = \boldsymbol{A}^{\mathrm{SI}}$. Formally, we want to prove that if $\boldsymbol{A}^k \leqslant \boldsymbol{A}^{\mathrm{SI}}$ and $\boldsymbol{A}^k \neq \boldsymbol{A}^{\mathrm{SI}}$ then $\boldsymbol{A}^{k+1} \leqslant \boldsymbol{A}^{\mathrm{SI}}$.

Indeed, assume that $\boldsymbol{A}^k \leqslant \boldsymbol{A}^{\mathrm{SI}}$ and $\boldsymbol{A}^k \neq \boldsymbol{A}^{\mathrm{SI}}$. Let $i \in N$ be the agent that is given a bundle in iteration $k+1$ of the mechanism. Since $\boldsymbol{A}^k \neq \boldsymbol{A}^{\mathrm{SI}}$, there exists an agent j who has a strictly smaller number of bundles in $\boldsymbol{A}^k$ than in $\boldsymbol{A}^{\mathrm{SI}}$. Note that $\boldsymbol{A}^{\mathrm{SI}}$ is a feasible allocation, hence the mechanism does have enough resources to allocate a bundle to j in iteration $k+1$, i.e., $j \in T_{k+1}$. It follows that

$$\mathrm{MaxDom}(\boldsymbol{A}^{k+1}) = \mathrm{MaxDom}(\boldsymbol{A}^k \uparrow i) \leqslant \mathrm{MaxDom}(\boldsymbol{A}^k \uparrow j) \leqslant \mathrm{MaxDom}(\boldsymbol{A}^{\mathrm{SI}}) \leqslant \frac{1}{n}.$$

Therefore, $\boldsymbol{A}^{k+1} \leqslant \boldsymbol{A}^{\mathrm{SI}}$. ■

The next lemma shows that the mechanism maintains EF1 at every step.

Lemma 4.3. SEQUENTIALMINMAX *satisfies EF1.*

Proof. For consistency, we say that the mechanism has an iteration 0 at the end of which the allocation is $\boldsymbol{A}^0 = \mathbf{0}$. Note that the allocation $\mathbf{0}$ is EF1. Suppose for contradiction that the allocation returned by the mechanism at the end is not EF1.

Iteration k_1: Let k_1 be the first iteration such that $\boldsymbol{A}^{k_1}$ is not EF1. Thus, $k_1 > 0$. Let $i \in N$ be the agent that is allocated a bundle in iteration k_1. There exists an agent $j \in N$ who envies i up to one bundle under $\boldsymbol{A}^{k_1}$, otherwise the allocation after iteration $k_1 - 1$ would not be EF1 as well, contradicting the fact that k_1 is the first such iteration. Since j envies i, it must hold that i has positive demand for every resource for which j has positive demand. That is,

$$\forall r \in R,\ b_{jr} > 0 \Rightarrow b_{ir} > 0. \tag{4.5}$$

Let x_i and x_j denote the number of bundles of i and j, respectively, in $\boldsymbol{A}^{k_1}$. It holds that $x_i \cdot \boldsymbol{b}_i \geqslant (x_j + 2) \cdot \boldsymbol{b}_j$, and therefore

$$\forall r \in R,\ x_i \cdot b_{ir} \geqslant (x_j + 2) \cdot b_{jr}. \tag{4.6}$$

For every $r \in R$ such that $b_{ir} > 0$, Equation (4.6) implies that $x_i \geqslant (x_j + 2) \cdot (b_{jr}/b_{ir})$. Let $\widehat{r} \in \arg\max_{DRF:r:b_{ir}>0}(b_{jr}/b_{ir})$. Then, we have

$$x_i \geqslant (x_j + 2) \cdot \frac{b_{j\widehat{r}}}{b_{i\widehat{r}}}. \tag{4.7}$$

Iteration k_2: Consider the iteration k_2 in which j was allocated its x_j-th bundle. Since i was allocated a bundle in iteration k_1, we have $k_2 \neq k_1$. Hence, $0 \leqslant k_2 < k_1$. Since allocations are monotonic, j has x_j bundles at the end of iteration k for every $k_2 \leqslant k \leqslant k_1$. If $k_2 = 0$, then j clearly does not envy i at the end of iteration k_2. If $k_2 > 0$ and j envies i at the end of iteration k_2, then j envied i up to one bundle at the end of iteration $k_2 - 1 < k_1$ since j had $x_j - 1$ bundles after iteration $k_2 - 1$. But this contradicts our assumption that k_1 is the first iteration such that A^{k_1} is not EF1. Hence, it must be the case that j does not envy i at the end of iteration k_2. However, j envies i (even up to one bundle) at the end of iteration k_1.

Iteration k_3: Take the smallest k_3 such that $k_2 < k_3 \leqslant k_1$ and j envies i at the end of iteration k_3. Clearly, the mechanism must have allocated one bundle to i in iteration k_3. Let x'_i be the number of bundles allocated to i at the end of iteration k_3 (hence $x'_i \geqslant 1$). Note that j must have exactly x_j bundles at the end of iteration k_3 since $k_2 < k_3 \leqslant k_1$. Further, since j does not envy i at the end of iteration $k_3 - 1$, there exists $r_1 \in R$ such that $(x'_i - 1) \cdot b_{ir_1} < (x_j + 1) \cdot b_{jr_1}$. First, this implies that $b_{jr_1} > 0$. Hence, we also have $b_{ir_1} > 0$ from Equation (4.5). Thus, we have $x'_i - 1 < (x_j + 1) \cdot b_{jr_1} / b_{ir_1}$. Since $\widehat{r} \in \arg\max_{DRF:r:b_{ir}>0} (b_{jr}/b_{ir})$, we have

$$x'_i - 1 < (x_j + 1) \cdot \frac{b_{j\widehat{r}}}{b_{i\widehat{r}}}. \tag{4.8}$$

Subtracting Equation (4.8) from Equation (4.7), we obtain $x_i - x'_i + 1 \geqslant b_{j\widehat{r}}/b_{i\widehat{r}}$. Thus, $x_i - x'_i + 1 \geqslant b_{jr}/b_{ir}$ for every $r \in R$ such that $b_{ir} > 0$. This implies that $(x_i - x'_i + 1) \cdot b_{ir} \geqslant b_{jr}$ for every $r \in R$ such that $b_{ir} > 0$. Moreover, if $b_{ir} = 0$ then we have $b_{jr} = 0$ from Equation (4.5), and $(x_i - x'_i + 1) \cdot b_{ir} = b_{jr} = 0$. Thus, for every $r \in R$, $(x_i - x'_i + 1) \cdot b_{ir} \geqslant b_{jr}$, i.e., $(x_i - x'_i + 1) \cdot \boldsymbol{b}_i \geqslant \boldsymbol{b}_j$. Note that at least $(x_i - x'_i + 1) \cdot \boldsymbol{b}_i$ resources were available at the beginning of iteration k_3 since the mechanism allocated $(x_i - x'_i + 1)$ bundles to i from iteration k_3 until k_1. In particular, the mechanism had enough resources to allocate a bundle to j at the beginning of iteration k_3, i.e., $j \in T_{k_3}$.

Next, since j envies i at the end of iteration k_3, we have

$$\forall r \in R,\ x'_i \cdot b_{ir} \geqslant (x_j + 1) \cdot b_{jr}. \tag{4.9}$$

Let r^*_i and r^*_j be dominant resources of agents i and j respectively. Then Equation (4.9) implies that

$$x'_i \cdot b_{ir^*_i} \geqslant x'_i \cdot b_{ir^*_j} \geqslant (x_j + 1) \cdot b_{jr^*_j}. \tag{4.10}$$

Since the mechanism allocated a bundle to i in iteration k_3, it must be the case that $i \in M_{k_3}$. Using our conclusion that $j \in T_{k_3}$, we have $x'_i \cdot b_{ir^*_i} \leqslant (x_j + 1) \cdot b_{jr^*_j}$. Therefore, Equation (4.10) holds with equalities. It follows that for every $r \in R$ such that $b_{ir} > 0$,

$$\frac{b_{jr}}{b_{ir}} \leqslant \frac{x'_i}{x_j + 1} = \frac{b_{jr^*_j}}{b_{ir^*_j}}, \tag{4.11}$$

where the first transition is due to Equation (4.9). Note that $b_{jr^*_j} > 0$, hence Equation (4.5) implies $b_{ir^*_j} > 0$. From Equation (4.11), we can see that $x'_i \cdot b_{ir^*_j} < (x_j + 2) \cdot b_{jr^*_j}$. Thus, j

does not envy i up to one bundle at the end of iteration k_3. This implies that $k_3 < k_1$. Since i was allocated a bundle in iteration k_1, we also have $x'_i < x_i$.

Iteration k_4: Consider the iteration k_4, $k_3 < k_4 \leqslant k_1$, in which i is allocated its $(x'_i + 1)$-th bundle. From Equation (4.6), we know that $x_i \cdot b_{ir^*_j} \geqslant (x_j + 2) \cdot b_{jr^*_j}$. Subtracting $x'_i \cdot b_{ir^*_j} = (x_j + 1) \cdot b_{jr^*_j}$ (equality holds due to Equation (4.11)), we get $(x_i - x'_i) \cdot b_{ir^*_j} \geqslant b_{jr^*_j}$. Again using Equation (4.11), we conclude that for every $r \in R$ such that $b_{ir} > 0$,

$$x_i - x'_i \geqslant \frac{b_{jr^*_j}}{b_{ir^*_j}} \geqslant \frac{b_{jr}}{b_{ir}}.$$

Thus, $(x_i - x'_i) \cdot b_{ir} \geqslant b_{jr}$ if $b_{ir} > 0$. Also, if $b_{ir} = 0$ then Equation (4.5) implies $b_{jr} = 0$ and we have $(x_i - x'_i) \cdot b_{ir} = b_{jr} = 0$. Hence, $(x_i - x'_i) \cdot \boldsymbol{b}_i \geqslant \boldsymbol{b}_j$. Note that at least $(x_i - x'_i) \cdot \boldsymbol{b}_i$ resources were available at the beginning of iteration k_4 since $x_i - x'_i$ bundles were allocated to i from iteration k_4 till k_1. Hence, there were enough resources to allocate a bundle to j at the beginning of iteration k_4, i.e., $j \in T_{k_4}$. Equation (4.10) implies that $(x_j + 1) \cdot b_{jr^*_j} < (x'_i + 1) \cdot b_{ir^*_i} \leqslant \mathrm{MaxDom}(A^{k_4})$. Note that this is still not a contradiction because allocating to j may not decrease $\mathrm{MaxDom}(A^{k_4})$.

Iteration k_5: Consider the first iteration k_5, $k_5 \leqslant k_4$, such that $\mathrm{MaxDom}(A^{k_5}) > (x_j + 1) \cdot b_{jr^*_j}$. Hence, we have

$$\mathrm{MaxDom}(A^{k_5 - 1}) \leqslant (x_j + 1) \cdot b_{jr^*_j}. \tag{4.12}$$

Let l be the agent that is allocated a bundle in iteration k_5. Thus, $A^{k_5} = A^{k_5-1} \uparrow l$. As j has at most x_j bundles at the end of iteration $k_5 - 1$, Equation (4.12) implies $\mathrm{MaxDom}(A^{k_5-1} \uparrow j) \leqslant (x_j + 1) \cdot b_{jr^*_j} < \mathrm{MaxDom}(A^{k_5}) = \mathrm{MaxDom}(A^{k_5-1} \uparrow l)$. Also, $j \in T_{k_4}$ and $k_5 \leqslant k_4$, hence $j \in T_{k_5}$. It follows that $l \notin M_{k_5}$, which is a contradiction since the mechanism allocated a bundle to l in iteration k_5. Thus, our original assumption that there exists an iteration where the allocation is not EF1 is false. ■

Crucially, as a special case of Theorem 4.5 we find that an allocation that is PO, SI, and EF1 always *exists* for indivisible tasks; this is not *a priori* clear, and in fact, we are not aware of a more direct proof of existence. Once we know such an allocation exists though, we can consider alternative methods of finding it. For example, the computation of such an allocation can be carried out using a mixed integer program (MIP). Using a MIP approach one can also optimize an objective subject to the three constraints; e.g., maximize the minimum dominant share to achieve the most egalitarian allocation among allocations that are PO, SI, and EF1.

Another potential approach is to use the properties of competitive equilibrium from equal incomes (CEEI). In the setting of Section 4.2 a competitive equilibrium exists and satisfies PO, SI, and EF [99]. While indivisibilities preclude the existence of a CEEI (obviously, as PO and EF are incompatible), conceivably it may be possible to prove the existence of an approximate version that is EF1 using an approach similar to Budish [44]. Even if such an approximate CEEI exists, its computation can be challenging [160].

In any case, as an algorithm for finding PO+SI+EF1 allocations, SEQUENTIALMIN-MAX has several major advantages over other potential approaches. First, since the mechanism allocates a bundle to some agent in each execution of the while loop, it can be seen that the mechanism terminates in time $\mathcal{O}(n^2 m/b^*)$, where $b^* = \min_{i \in N} \max_{r \in R} b_{ir}$; i.e., $1/b^*$ is an upper bound on the number of bundles that can be allocated to an agent. Second, it can be implemented *dynamically*. Indeed, even in a realistic setting where agents change their demands over time, or arrive and depart, we can still carry out the policy of allocating to an agent minimizing the maximum dominant share. Specifically, at the beginning of iteration t we can compute the sets T_k and M_k regardless of the currently present agents and their current demands.

4.6 Related Work

The work of Ghodsi et al. [99] has quickly attracted significant attention from the algorithmic economics community. Consequently, during the process of preparing the conference version of this work, several independent (from our work and each other) related papers have become available. While none of these papers study social welfare maximization (our Section 4.4) or indivisibilities (our Section 4.5), there is a partial overlap with our results on extending DRF (Section 4.3). In the following, we discuss each of these papers in detail, starting from the original DRF paper.

Ghodsi et al. [99] focus on the theoretical comparison of DRF to other potential mechanisms, and provide a systematic empirical study of the computational properties of DRF. Ghodsi et al. also mention that one can extend DRF to take into account weighted agents. However, the discussion of this extension is informal, and it is unclear whether DRF maintains its properties. In addition, although Ghodsi et al. do not assume that agents have strictly positive demands for every resource, this issue is addressed rather informally. Specifically, their proof of SP relies on a loose simulation argument that, we believe, may be insufficient when some agents do not demand every resource. We introduce an alternative technical framework that allows us (in Section 4.3) to rigorously establish stronger properties (GSP instead of SP) while simultaneously tackling agent weights and zero demands on some resources.

A manuscript by Li and Xue [131] provides a characterization of mechanisms that satisfy desirable properties under Leontief preferences. As a consequence of their results, one obtains a formal proof of GSP for certain mechanisms, even when agents are weighted. Curiously, their results do not capture DRF itself due to a technical assumption. More importantly, Li and Xue assume strictly positive demands; we argue in Section 4.3 that the case of zero demands is important, and our results suggest that it is also technically challenging.

Another manuscript by Friedman et al. [97] explores the relations between resource allocation under Leontief preferences and bargaining theory. Among other results, they introduce a family of *weighted* DRF mechanisms. It is important to recognize that here the weights are not exogenously given, but are a way to induce different variations of DRF as mechanisms for allocating to *unweighted* agents. In addition, Friedman et al.

explicitly assume strictly positive demands for simplicity. Under this assumption, they show that every weighted DRF mechanism satisfying additional technical properties is GSP.

We mention two additional related papers that are disjoint from ours in terms of results. Dolev et al. [78] also study resource allocation under Leontief preferences. They consider an alternative fairness criterion, which they call *no justified complaints*. They compare allocations that satisfy this criterion with DRF allocations. Their main result is the existence of allocations that satisfy no justified complaints. Gutman and Nisan [108] present polynomial time algorithms for computing allocations under a family of mechanisms that includes DRF. Moreover, they show that a competitive equilibrium (discussed in passing below) achieves the notion of fairness of Dolev et al. [78], thereby leading to a polynomial time algorithm for computing allocations with this property. The two new mechanisms that we present are both polynomial time.

4.7 Discussion

This chapter enhances our understanding of resource allocation in settings with multiple resources. First, we assumed that tasks are divisible and showed that EXTENDED-DRF is PO, SI, EF, and GSP even when the agents have initial endowments and may not demand every resource. Second, we observed that no SI or EF mechanism can provide an approximation ratio smaller than m for the social welfare, and proved that the same property holds when one requires SP. Third, under indivisible tasks, we showed that SEQUENTIALMINMAX is PO, SI, and EF1. Nevertheless, several key challenges remain unresolved.

Our model of Section 4.5 relaxes the assumption that tasks are divisible, but we do make the assumption that there is one pool of divisible resources. As discussed in passing by Ghodsi et al. [99], in practice clusters consist of many small machines. This issue of *resource fragmentation* complicates matters further, and requires careful attention.

In addition, we spent some time studying a setting where each agent has multiple divisible tasks. We focused on a specific utility function: an agent's utility is the sum of dominant shares of its tasks, under the optimal division of its allocated resources between its tasks. Such a utility function can be motivated by routing settings where resources are links and tasks correspond to different paths from a source to a sink. It is possible to show that a competitive equilibrium exists in this setting, and the resulting allocation satisfies PO, SI, and EF. However, we were unable to settle the existence of a mechanism satisfying PO, SI, EF, and SP, even under restricted utility functions.

Finally, our analysis is restricted to a static setting where demands are revealed and a single allocation is made. Above we touched on the issue of dynamic implementation, and indeed in practice agents may arrive and depart, or change their demands over time (cf. [163, 212] for analogous settings with money). What theoretical guarantees would we look for in such a setting? In particular, what would be the appropriate notion of fairness in a dynamic resource allocation setting? Some answers are in Chapter 5, but many fundamental questions remain open.

Chapter 5

Dynamic Fair Division

5.1 Introduction

In the previous chapters, we studied a variety of fair division settings involving additive utilities, dichotomous utilities, and Leontief preferences. In fact, the fair division literature offers a wide spectrum of other possibilities that apply to different real-world domains. Nevertheless, some aspects of such domains are beyond the current scope of fair division theory. Perhaps most importantly, the literature fails to capture the *dynamics* of the system. In other words, for centuries fair division theory has addressed the static problem of finding a one-shot fair allocation of a known set of resources among a known set of agents. In many real-world systems, it is typically not the case that all the agents are present in the system at any given time; agents may arrive and depart, and the system must be able to adjust the allocation of resources. Even on the conceptual level, dynamic settings challenge some of the premises of fair division theory. For example, if one agent arrives before another, the first agent should intuitively have priority; what does fairness mean in this context?

Focusing on the setting from the previous chapter involving the allocation of computational resources in computing systems, we introduce the concepts necessary to answer this question, and design novel mechanisms that satisfy our proposed desiderata. Our contribution is therefore twofold: we *design more realistic resource allocation mechanisms for multiagent systems* that provide theoretical guarantees, and at the same time we *expand the scope of fair division theory* to capture dynamic settings.

5.2 Dynamic Resource Allocation: A New Model

Let us first briefly review the static model studied in the previous chapter. In this model, a set of agents $N = \{1, \ldots, n\}$ have Leontief preferences over a set of resources R (where $|R| = m$). We represent the (normalized) demand vector of agent i by $\boldsymbol{d}_i = \langle d_{i1}, \ldots, d_{im} \rangle$, and the set of all possible normalized demand vectors by $\mathcal{D}$. Let $\boldsymbol{d}_{\leqslant k} = \langle \boldsymbol{d}_1, \ldots, \boldsymbol{d}_k \rangle$ denote the demand vectors of agents 1 through k. Similarly, let $\boldsymbol{d}_{>k} = \langle \boldsymbol{d}_{k+1}, \ldots, \boldsymbol{d}_n \rangle$ denote the demand vectors of agents $k+1$ through n.

Recall that an allocation $\boldsymbol{A}$ allocates a fraction A_{ir} of resource r to agent i, subject to the feasibility condition $\sum_{i \in N} A_{ir} \leqslant 1$ for all $r \in R$, and that the utility of agent i under this allocation is given by

$$u_i(\boldsymbol{A}_i) = \max\{y \in \mathbb{R}_+ : \forall r \in R,\ A_{ir} \geqslant y \cdot d_{ir}\}.$$

An allocation $\boldsymbol{A}$ is called *non-wasteful* if for every agent i there exists $y \in \mathbb{R}_+$ such that for all $r \in R$, $A_{ir} = y \cdot d_{ir}$. For a non-wasteful allocation, the utility of an agent is the share of its dominant resource allocated to the agent. If $\boldsymbol{A}$ is a non-wasteful allocation then for all $i \in N$,

$$u_i(\boldsymbol{A}'_i) > u_i(\boldsymbol{A}_i) \Rightarrow \forall r \in R,\ A'_{ir} > A_{ir}. \tag{5.1}$$

For allocations $\boldsymbol{A}$ over agents in $S \subseteq N$ and $\boldsymbol{A}'$ over agents in $T \subseteq N$ such that $S \subseteq T$, we say that $\boldsymbol{A}'$ is an *extension* of $\boldsymbol{A}$ to T if $A'_{ir} \geqslant A_{ir}$ for every agent $i \in S$ and every resource r. When $S = T$, we simply say that $\boldsymbol{A}'$ is an extension of $\boldsymbol{A}$.

We now introduce our dynamic resource allocation model. In our model, agents arrive at different times and do not depart (see Section 5.7 for a discussion of this point). We assume that agent 1 arrives first, then agent 2, and in general agent k arrives after agents $1, \ldots, k-1$; we say that agent k arrives in *step* k. An agent reports its demand when it arrives and the demand does not change over time. Thus, at step k, demand vectors $\boldsymbol{d}_{\leqslant k}$ are known, and demand vectors $\boldsymbol{d}_{>k}$ are unknown. A *dynamic resource allocation mechanism* operates as follows. At each step k, the mechanism takes as input the reported demand vectors $\boldsymbol{d}_{\leqslant k}$ and outputs an allocation $\boldsymbol{A}^k$ over the agents present in the system. Crucially, we assume that allocations are *irrevocable*, i.e., $A^k_{ir} \geqslant A^{k-1}_{ir}$ for every step $k \geqslant 2$, every agent $i \leqslant k-1$, and every resource r. We also assume that the mechanism knows the total number of agents n in advance.

Irrevocability can be justified in various settings, e.g., in cases where resources are committed to long-term projects. One example is that of a research cluster shared between faculty members at a university. In such a cluster, the total number of faculty members who can access the cluster (denoted n in our setting) is known to the mechanism in advance — as we assume in our model.

Previous work on static resource allocation [see, e.g., 99, 165] focused on designing mechanisms that satisfy four prominent desiderata. Three of these — two fairness properties and one game-theoretic property — immediately extend to the dynamic setting.

1. *Sharing Incentives (SI).* We say that a dynamic allocation mechanism satisfies SI if $u_i(\boldsymbol{A}^k_i) \geqslant u_i(\langle 1/n, \ldots, 1/n \rangle)$ for all steps k and all agents $i \leqslant k$. In words, when an agent arrives it receives an allocation that it likes at least as much as an equal split of the resources. This models a setting where agents have made equal contributions to the system and hence have equal entitlements. In such cases, the contributions are typically recorded, which allows the mechanism to know the total number of agents n in advance, as assumed in our setting.
2. *Envy Freeness (EF).* A dynamic allocation mechanism is EF if $u_i(\boldsymbol{A}^k_i) \geqslant u_i(\boldsymbol{A}^k_j)$ for all steps k and all agents $i, j \leqslant k$, that is, an agent that is present would never prefer the allocation of another agent.

3. *Strategyproofness (SP)*. A dynamic allocation mechanism is SP if no agent can misreport its demand vector and be strictly better off at any step k, regardless of the reported demands of other agents. Formally, a dynamic allocation mechanism is SP if for any agent $i \in N$ and any step k, if $\boldsymbol{A}_i^k$ is the allocation to agent i at step k when agent i reports its true demand vector and $\boldsymbol{B}_i^k$ is the allocation to agent i at step k when agent i reports a different demand vector (in both cases all the other agents report their true demand vectors), then $u_i(\boldsymbol{A}_i^k) \geqslant u_i(\boldsymbol{B}_i^k)$. We avoid introducing additional notations that will not be required later.

In the static setting, the fourth prominent axiom, *Pareto optimality (PO)*, means that the mechanism's allocation is not Pareto dominated by any other allocation. Of course, in the dynamic setting it is unreasonable to expect the allocation in early stages to be Pareto undominated, because we need to save resources for future arrivals (recall that allocations are irrevocable). We believe though that the following definition naturally extends PO to our dynamic setting.

4. *Dynamic Pareto Optimality (DPO).* A dynamic allocation mechanism is DPO if at each step k, the allocation $\boldsymbol{A}^k$ returned by the mechanism is not Pareto dominated by any other allocation $\boldsymbol{B}^k$ that allocates up to a (k/n)-fraction of each resource among the k agents present in the system. Put another way, at each step the allocation should not be Pareto dominated by any other allocation that only redistributes the collective entitlements of the agents present in the system among those agents.

It is straightforward to verify that a non-wasteful mechanism (a mechanism returning a non-wasteful allocation at each step) satisfies DPO if and only if the allocation returned by the mechanism at each step k uses at least a (k/n)-fraction of at least one resource (the assumption of strictly positive demands plays a role here).

Before moving on to possibility and impossibility results, we give examples that illustrate how various combinations of the properties constrain the allocation of resources.

Example 5.1 (Satisfying Sharing Incentives (SI) and Dynamic Pareto Optimality (DPO))**.** In this chapter, we only consider non-wasteful allocations. Hence, as described above, DPO is equivalent to allocating at least a (k/n)-fraction of at least one resource in every step k, when allocations are proportional. On the other hand, if a mechanism seeks to satisfy SI, it cannot allocate more than a (k/n)-fraction of any resource in step k. Indeed, if more than a (k/n)-fraction of resource r is allocated at step k, and every agent arriving after step k reports r as its dominant resource, the mechanism would not have enough of resource r left to allocate each of them at least a $(1/n)$-fraction of r, as required by SI. Thus, a non-wasteful mechanism satisfying both SI and DPO must allocate, in every step k, exactly a (k/n)-fraction of some resource and at most a (k/n)-fraction of every other resource. In other words, in every step k, the mechanism has a pool of available resources — which contains a (k/n)-fraction of each resource, minus the fraction already allocated — to be allocated to the k agents that are currently present. The mechanism can only allocate from this pool, and must exhaust at least one resource from the pool.

Example 5.2 (Understanding Strategyproofness (SP))**.** In this example, we take a mechanism that may seem SP at first glance, and show that it violates our definition of SP. For simplicity, we will allow the agents to have possibly zero demands for some of the resources in this example. This allows beneficial manipulations for the following simple mechanism, which we call DYNAMIC DICTATORSHIP. (We note that DYNAMIC DICTATORSHIP is otherwise strategyproof for strictly positive demands — see the discussion following Theorem 5.1.) At each step k, the mechanism allocates a $1/n$ share of each resource to agent k, takes back the shares of different resources that the agent cannot use, and then allocates resources to the k present agents in the order of their arrival using serial dictatorship, that is, it allocates to each agent as many resources as the agent can use, and then proceeds to the next agent. The mechanism keeps allocating until a k/n share of at least one resource is allocated. Note that the mechanism trivially satisfies SI because it allocates resources as valuable as an equal split to each agent as soon as it arrives. The mechanism would satisfy DPO in our standard setting with non-zero demands, because it is non-wasteful and at every step k it allocates a k/n fraction of at least one resource. Intuitively, it seems that the first agent should not gain by reporting a false demand vector because in each round it gets to pick first and is allowed to take as much as it can use from the available pool of resources. We show that this intuition is incorrect. Let us denote the pool of resources available to the mechanism in any step by a vector of the fraction of each available resource. Consider the case of four agents (agents 1, 2, 3, and 4), and three resources (R_1, R_2, and R_3). Let the true demand vectors of the agents be as follows:

$$d_1 = \langle 1, 0.5, 0.5 \rangle,\ d_2 = \langle 0, 1, 1 \rangle,\ d_3 = \langle 1, 0.5, 0 \rangle,\ d_4 = \langle 0, 1, 0.5 \rangle.$$

Figure 5.1 shows the allocations returned by DYNAMIC DICTATORSHIP in various steps when all agents report their true demand vectors. Now, suppose agent 1 raises its demand for R_3 by reporting a false demand vector $\langle 1, 0.5, 1 \rangle$. In this case the allocations returned by the mechanism in various steps are shown in Figure 5.2. We can see that the manipulation makes agent 1 strictly worse off in step 2, but strictly better off in the final step. Our definition of SP requires that an agent should not be able to benefit in any step of the process by lying — thus DYNAMIC DICTATORSHIP is not SP.

5.2.1 Impossibility Result

Ideally, we would like to design a dynamic allocation mechanism that is SI, EF, SP, and DPO. However, we show that even satisfying EF and DPO simultaneously is impossible.

Theorem 5.1. *Let $n \geqslant 3$ and $m \geqslant 2$. Then no dynamic resource allocation mechanism satisfies EF and DPO.*

Proof. Consider a setting with three agents and two resources. Agents 1 and 2 have demand vectors $\langle 1, 1/9 \rangle$ and $\langle 1/9, 1 \rangle$, respectively (i.e., $d_{11} = 1$, $d_{12} = 1/9$, etc.). At step 2 (after the second agent arrives), at least one of the two agents must be allocated at

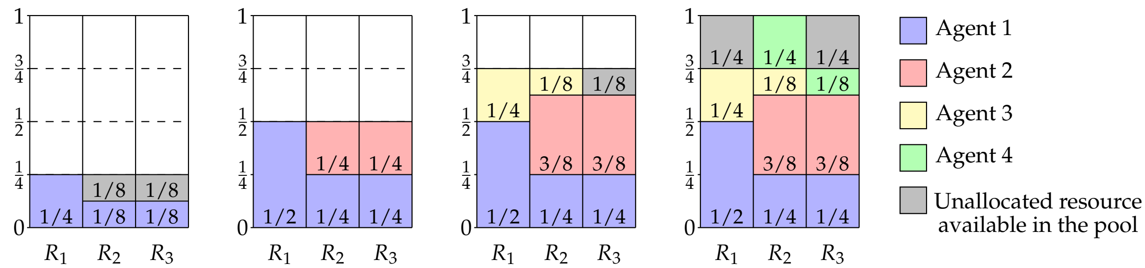

(a) Step 1 (b) Step 2 (c) Step 3 (d) Step 4

Figure 5.1: Allocations returned by DYNAMIC DICTATORSHIP when agent 1 reports its true demand vector.

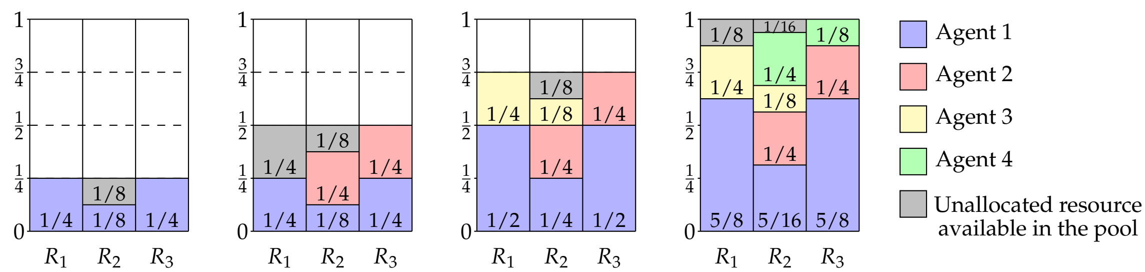

(a) Step 1 (b) Step 2 (c) Step 3 (d) Step 4

Figure 5.2: Allocations returned by DYNAMIC DICTATORSHIP when agent 1 manipulates.

least an $x = 3/5$ share of its dominant resource. Suppose for contradiction that the two agents are allocated x' and x'' shares of their dominant resources where $0 < x', x'' < x$. Then, the total fractions of the two resources allocated at step 2 would be $x' + x'' \cdot (1/9)$ and $x'' + x' \cdot (1/9)$, both less than $x + x \cdot (1/9) = 2/3$, violating DPO. Without loss of generality, assume that agent 1 is allocated at least an $x = 3/5$ share of its dominant resource (resource 1) at step 2. If agent 3 reports the demand vector $\langle 1, 1/9 \rangle$ — identical to that of agent 1 — then it can be allocated at most a $2/5$ share of its dominant resource (resource 1), and would envy agent 1.

It is easy to extend this argument to the case of $n > 3$, by adding $n - 3$ agents with demand vectors that are identical to the demand vector of agent 3. Once again, it can be verified that at the end of step 2, at least one of the first two agents (w.l.o.g., agent 1) must be allocated at least a $9/(5n)$ share of its dominant resource. If we take the remaining resources (in particular, at most a $1 - 9/(5n)$ share of resource 1), and divide them among the remaining $n - 2$ agents that have demand vectors identical to that of agent 1, at least one of them will get at most a $(1 - 9/(5n))/(n - 2) < 9/(5n)$ share of its dominant resource, and will envy agent 1. To extend to the case of $m > 2$, let all agents have negligibly small demands for the additional resources. ■ (Proof of Theorem 5.1)

It is interesting to note that if either EF or DPO is dropped, the remaining three axioms can be easily satisfied. For example, the trivial mechanism EQUAL SPLIT that just gives every agent a $1/n$ share of each resource when it arrives satisfies SI, EF and SP. Achieving SI, DPO, and SP is also simple. Indeed, consider the DYNAMIC DICTATORSHIP mechanism from Example 5.2. The example explains why DYNAMIC DICTATORSHIP satisfies both SI and DPO. Even though DYNAMIC DICTATORSHIP is not SP under possibly zero demands (as shown in the example), it is clearly SP for strictly positive demands (as assumed throughout this chapter). When agent k arrives in step k, it is allocated a $1/n$ share of its dominant resource (and other resources in proportion), and subsequently agent 1 is allocated resources until a k/n share of at least one resource is exhausted. Since every agent requires the exhausted resource due to strictly positive demands, the allocation stops. In summary, all agents except agent 1 receive exactly a $1/n$ share of their dominant resource when they arrive, and do not receive any resources later on; hence, they cannot gain by reporting a false demand vector. In step k, agent 1 receives as much resources as it can from the pool of resources that remain after allocating to agents 2 through k a $1/n$ share of their dominant resource from an original pool that contains a k/n share of each resource. Therefore, agent 1 also cannot gain from manipulation.

While both EQUAL SPLIT and DYNAMIC DICTATORSHIP satisfy maximal subsets of our proposed desiderata, neither is a compelling mechanism. Since these mechanisms are permitted by dropping EF or DPO entirely, we instead explore relaxations of EF and DPO that rule these mechanisms out and guide us towards more compelling mechanisms.

5.3 Relaxing Envy Freeness

Recall that DPO requires a mechanism to allocate at least a k/n fraction of at least one resource at step k, for every $k \in \{1, \ldots, n\}$. Thus the mechanism sometimes needs to allocate a large amount of resources to agents arriving early, potentially making it impossible for the mechanism to prevent the late agents from envying the early agents. In other words, when an agent i enters the system it may envy some agent j that arrived before i did; this is inevitable in order to be able to satisfy DPO. However, it would be unfair to agent i if agent j were allocated more resources since agent i arrived while i still envied j. To distill this intuition, we introduce the following dynamic version of EF.

2′. *Dynamic Envy Freeness (DEF)*. A dynamic allocation mechanism is DEF if at any step an agent i envies an agent j *only if* j arrived before i did *and* j has not been allocated any resources since i arrived. Formally, for every $k \in \{1, \ldots, n\}$, if $u_i(\boldsymbol{A}_j^k) > u_i(\boldsymbol{A}_i^k)$ then $j < i$ and $\boldsymbol{A}_j^k = \boldsymbol{A}_j^{i-1}$.

Walsh [198] studied a dynamic cake cutting setting and proposed *forward EF*, which requires that an agent not envy any agent that arrived later. This notion is weaker than DEF because it does not rule out the case where an agent i envies an agent j that arrived earlier *and* j received resources since i arrived. In our setting, even the trivial mechanism DYNAMIC DICTATORSHIP (see Section 5.2.1) satisfies forward EF, but fails to satisfy our stronger notion of DEF.

We next construct a dynamic resource allocation mechanism — DYNAMIC DRF — that achieves the relaxed fairness notion of DEF, together with SI, DPO, and SP. The mechanism is given as Algorithm 6.

ALGORITHM 6: DYNAMIC DRF

Data: Demands $\boldsymbol{d}$

Result: Allocation A^k at each step k

$k \leftarrow 1$;

while $k \leqslant n$ **do**

  $\{x_i^k\}_{i=1}^k \leftarrow$ Solution of the LP in the box below;

  $A_{ir}^k \leftarrow x_i^k \cdot d_{ir}, \forall i \leqslant k$;

  $k \leftarrow k + 1$;

end

Maximize M^k
subject to
$x_i^k \geqslant M^k, \forall i \leqslant k$
$x_i^k \geqslant x_i^{k-1}, \forall i \leqslant k-1$
$\sum_{i=1}^k x_i^k \cdot d_{ir} \leqslant k/n, \forall r \in R$

Intuitively, at each step k the mechanism starts from the current allocation among the present agents and keeps allocating resources to agents that have the minimum dominant share *at the same rate*, until a k/n fraction of at least one resource is allocated. Always allocating to agents that have the minimum dominant share ensures that agents are not allocated any resources while they are envied. This *water-filling* mechanism is a dynamic adaptation of the dominant resource fairness (DRF) mechanism proposed by Ghodsi et al. [99]. See Figure 5.3 for an example.

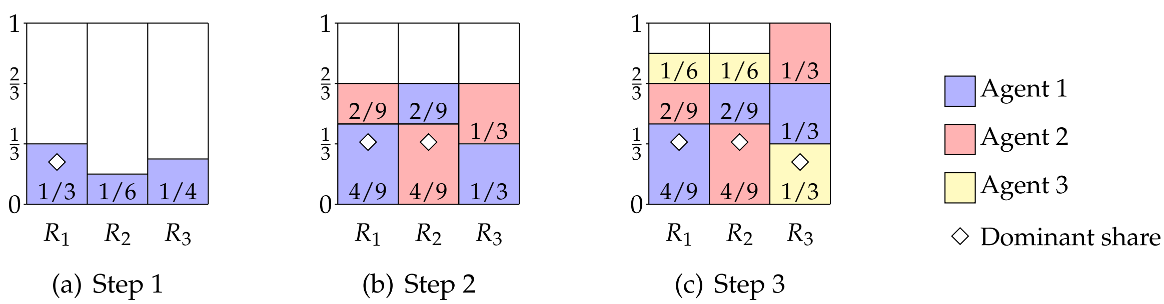

Figure 5.3: Allocations returned by DYNAMIC DRF at various steps for 3 agents with demands $\boldsymbol{d}_1 = \langle 1, 1/2, 3/4 \rangle$, $\boldsymbol{d}_2 = \langle 1/2, 1, 3/4 \rangle$, and $\boldsymbol{d}_3 = \langle 1/2, 1/2, 1 \rangle$, and three resources R_1, R_2, and R_3. Agent 1 receives a $1/3$ share of its dominant resource at step 1. At step 2, water-filling drives the dominant shares of agents 1 and 2 up to $4/9$. At step 3, however, agent 3 can only receive a $1/3$ dominant share and the allocations of agents 1 and 2 remain unchanged.

Theorem 5.2. DYNAMIC DRF *satisfies SI, DEF, DPO, and SP, and can be implemented in polynomial time.*

Proof. First we show that DYNAMIC DRF satisfies SI. We need to prove that $x_i^k \geqslant 1/n$ for all agents $i \leqslant k$ at every step $k \in \{1, \ldots, n\}$. We prove this by induction on k. For the base case $k = 1$, it is easy to see that $x_1^1 = 1/n$ and $M^1 = 1/n$ is a solution of the LP of DYNAMIC DRF and hence the optimal solution satisfies $x_1^1 \geqslant M^1 \geqslant 1/n$ (in fact, there is an equality). Assume that this is true at step $k-1$ and let us prove the claim for step k, where $k \in \{2, \ldots, n\}$. At step k, one feasible solution of the LP is given by $x_i^k = x_i^{k-1}$ for agents $i \leqslant k-1$, $x_k^k = 1/n$ and $M^k = 1/n$. To see this, note that it trivially satisfies the first two constraints of the LP, because by the induction hypothesis we have $x_i^{k-1} \geqslant 1/n$ for $i \leqslant k-1$. Furthermore, in the proposed feasible solution, for any $r \in R$ we have

$$\sum_{i=1}^{k} x_i^k \cdot d_{ir} = \sum_{i=1}^{k-1} x_i^{k-1} \cdot d_{ir} + \frac{1}{n} \cdot d_{kr} \leqslant \frac{k-1}{n} + \frac{1}{n} \leqslant \frac{k}{n},$$

where the first transition follows from the construction of the feasible solution and the second transition holds because $\{x_i^{k-1}\}_{i=1}^{k-1}$ satisfies the LP of step $k-1$, and in particular the third constraint of the LP. Since a feasible solution achieves $M^k = 1/n$, the optimal solution achieves $M^k \geqslant 1/n$. Thus in the optimal solution $x_i^k \geqslant M^k \geqslant 1/n$ for all $i \leqslant k$, which is the requirement for SI.

Next we show that DYNAMIC DRF satisfies DPO. Observe that at any step k, the third constraint of the LP must be tight for at least one resource in the optimal solution (otherwise every x_i^k along with M^k can be increased by a sufficiently small quantity, contradicting the optimality of M^k). Thus, at each step k the (non-wasteful) mechanism allocates a k/n fraction of at least one resource, which implies that the mechanism satisfies DPO.

To prove that the mechanism satisfies DEF and SP, we first prove several useful lemmas about the allocations returned by the mechanism. In the proof below, M^k and x_i^k refer to the *optimal* solution of the LP in step k. Furthermore, we assume that $x_i^k = 0$ for agents $i > k$ (i.e., agents not present in the system are not allocated any resources). We begin with the following lemma, which essentially shows that if an agent is allocated some resources in a step using water-filling, then the agent's dominant share after the step will be the minimum among the present agents.

Lemma 5.1. *At every step $k \in \{1, \ldots, n\}$, it holds that $x_i^k = \max(M^k, x_i^{k-1})$ for all agents $i \leqslant k$.*

Proof. Consider any step $k \in \{1, \ldots, n\}$. From the first and the second constraints of the LP it is evident that $x_i^k \geqslant M^k$ and $x_i^k \geqslant x_i^{k-1}$ (note that $x_k^{k-1} = 0$), thus $x_i^k \geqslant \max(M^k, x_i^{k-1})$ for all $i \leqslant k$. Suppose for contradiction that $x_i^k > \max(M^k, x_i^{k-1})$ for some $i \leqslant k$. Then x_i^k can be reduced by a sufficiently small $\varepsilon > 0$ without violating any constraints. This makes the third constraint of the LP loose by at least $\varepsilon \cdot d_{ir}$, for every resource $r \in R$. Consequently, the values of x_j^k for $j \neq i$ and M^k can be increased by a sufficiently small $\delta > 0$ without violating the third constraint of the LP. Finally, ε (and correspondingly δ) can be chosen to be small enough so that $x_i^k \geqslant M^k$ is not violated. It follows that the value of M^k can be increased, contradicting the optimality of M^k. ■ (Proof of Lemma 5.1)

Next we show that at each step k, the dominant shares of agents 1 through k are monotonically non-increasing with their time of arrival. This is intuitive because at every step k, agent k enters with zero dominant share and subsequently we perform water-filling, hence monotonicity is preserved.

Lemma 5.2. *For all agents $i, j \in N$ such that $i < j$, we have $x_i^k \geqslant x_j^k$ at every step $k \in \{1, \dots, n\}$.*

Proof. Fix any two agents $i, j \in N$ such that $i < j$. We prove the lemma by induction on k. The result trivially holds for $k < j$ since $x_j^k = 0$. Assume that $x_i^{k-1} \geqslant x_j^{k-1}$ where $k \in \{j, \dots, n\}$. At step k, we have $x_i^k = \max(M^k, x_i^{k-1}) \geqslant \max(M^k, x_j^{k-1}) = x_j^k$, where the first and the last transition follow from Lemma 5.1 and the second transition follows from our induction hypothesis. ■ (Proof of Lemma 5.2)

The following lemma shows that if agent j has a greater dominant share than agent i at some step, then j must have arrived before i and j must not have been allocated any resources since i arrived. Observe that this is very close to the requirement of DEF.

Lemma 5.3. *At any step $k \in \{1, \dots, n\}$, if $x_j^k > x_i^k$ for some agents $i, j \leqslant k$, then $j < i$ and $x_j^k = x_j^{i-1}$.*

Proof. First, note that $j < i$ trivially follows from Lemma 5.2. Suppose for contradiction that $x_j^k > x_j^{i-1}$ (it cannot be smaller because allocations are irrevocable). Then there exists a step $t \in \{i, \dots, k\}$ such that $x_j^t > x_j^{t-1}$. Now Lemma 5.1 implies that $x_j^t = M^t \leqslant x_i^t$, where the last transition follows because x_i^t satisfies the second constraint of the LP at step t (note that $i \leqslant t$). However, $x_j^t \geqslant x_i^t$ due to Lemma 5.2. Thus, $x_j^t = x_i^t$. Now using Lemma 5.1, $x_j^{t+1} = \max(M^{t+1}, x_j^t) = \max(M^{t+1}, x_i^t) = x_i^{t+1}$. Extending this argument using a simple induction shows that $x_j^{t'} = x_i^{t'}$ for every step $t' \geqslant t$, in particular, $x_j^k = x_i^k$, contradicting our assumption. ■ (Proof of Lemma 5.3)

We proceed to show that DYNAMIC DRF satisfies DEF. We need to prove that for any step $k \in \{1, \dots, n\}$ and any agents $i, j \leqslant k$, if agent i envies agent j in step k (i.e., $u_i(\boldsymbol{A}_j^k) > u_i(\boldsymbol{A}_i^k)$), then $j < i$ and $x_j^k = x_j^{i-1}$. First, note that $u_i(\boldsymbol{A}_j^k) > u_i(\boldsymbol{A}_i^k)$ trivially implies that $x_j^k > x_i^k$, otherwise for the dominant resource r_i^* of agent i, we would have $A_{ir_i^*}^k = x_i^k \geqslant x_j^k \geqslant x_j^k \cdot d_{jr_i^*} = A_{jr_i^*}^k$ and agent i would not envy agent j. Now DEF follows from Lemma 5.3.

To prove that DYNAMIC DRF is SP, suppose for contradiction that an agent $i \in N$ can report an untruthful demand vector $\boldsymbol{d}_i'$ such that the agent is strictly better off in at least one step. Let k be the first such step. Denote by $\widehat{x}_j^t$ the dominant share of an agent j at step t with manipulation (for agent i, this is the share of the dominant resource of the untruthful demand vector) and similarly, denote by $\widehat{M}^t$ the value of M^t in the optimal solution of the LP of step t with manipulation.

Lemma 5.4. *$\widehat{x}_j^k \geqslant x_j^k$ for every agent $j \leqslant k$.*

Proof. For any agent j such that $x_j^k > x_i^k$, we have

$$x_j^k = x_j^{i-1} = \widehat{x}_j^{i-1} \leqslant \widehat{x}_j^k.$$

Here, the first transition follows from Lemma 5.3, the second transition holds because manipulation by agent i does not affect the allocation at step $i-1$, and the third transition follows from the LP. For any agent j with $x_j^k \leqslant x_i^k$, we have

$$x_j^k \leqslant x_i^k < \widehat{x}_i^k = \widehat{M}^k \leqslant \widehat{x}_j^k.$$

The second transition is true because if $\widehat{x}_i^k \leqslant x_i^k$ then agent i could not be better off as the true dominant share it receives with manipulation would be no more than it received without manipulation. To justify the third transition, note that agent i must be allocated some resources at step k with manipulation. If $k = i$, this is trivial, and if $k > i$, this follows because otherwise k would not be the first step when agent i is strictly better off as we would have $u_i(\widehat{\boldsymbol{A}}_i^{k-1}) = u_i(\widehat{\boldsymbol{A}}_i^k) > u_i(\boldsymbol{A}_i^k) \geqslant u_i(\boldsymbol{A}_i^{k-1})$, where $\widehat{\boldsymbol{A}}_i^k$ denotes the allocation to agent i at step k with manipulation. Thus, $\widehat{x}_i^k > \widehat{x}_i^{k-1}$, and the third transition now follows from Lemma 5.1. The last transition holds because $\widehat{x}_j^k$ satisfies the first constraint of the LP of step k. Thus, we conclude that $\widehat{x}_j^k \geqslant x_j^k$ for all agents $j \leqslant k$. ■ (Proof of Lemma 5.4)

Now, the mechanism satisfies DPO and thus allocates at least a k/n fraction of at least one resource at step k without manipulation. Let r be such a resource. Then the fraction of resource r allocated at step k with manipulation is

$$\widehat{x}_i^k \cdot d'_{ir} + \sum_{\substack{j \leqslant k \\ s.t. j \neq i}} \widehat{x}_j^k \cdot d_{jr} > x_i^k \cdot d_{ir} + \sum_{\substack{j \leqslant k \\ s.t. j \neq i}} x_j^k \cdot d_{jr} \geqslant k/n.$$

To justify the inequality, note that $\widehat{x}_i^k \cdot d'_{ir} > x_i^k \cdot d_{ir}$ by Equation (5.1) (as agent i is strictly better off), and in addition $\widehat{x}_j^k \geqslant x_j^k$ for every $j \leqslant k$. However, this shows that more than a k/n fraction of resource r must be allocated at step k with manipulation, which is impossible due to the third constraint of the LP. Hence, a successful manipulation is impossible, that is, DYNAMIC DRF is SP.

Finally, note that the LP has a linear number of variables and constraints, therefore the mechanism can be implemented in polynomial time. ■ (Proof of Theorem 5.2)

5.4 Relaxing Dynamic Pareto Optimality

We saw (Theorem 5.1) that satisfying EF and DPO is impossible. We then explored an intuitive relaxation of EF. Despite the positive result (Theorem 5.2), the idea of achieving absolute fairness — as conceptualized by EF — in our dynamic setting is compelling.

As a straw man, consider waiting for all the agents to arrive and then using any EF static allocation mechanism. However, this scheme is highly inefficient, e.g., it is easy to see that one can always allocate each agent at least a $1/n$ share of its dominant resource (and other resources in proportion) as soon as it arrives and still maintain EF at every step. How much more can be allocated at each step? We put forward a general answer to this question using a relaxed notion of DPO that requires a mechanism to allocate as many resources as possible while ensuring that EF can be achieved in the future, but first we require the following definition. Given a step $k \in \{1, \dots, n\}$, define an allocation $\boldsymbol{A}$ over the k present agents with demands $\boldsymbol{d}_{\leqslant k}$ to be *EF-extensible* if it can be extended to an EF allocation over all n agents with demands $\boldsymbol{d} = (\boldsymbol{d}_{\leqslant k}, \boldsymbol{d}_{>k})$, for all possible future demand vectors $\boldsymbol{d}_{>k} \in \mathcal{D}^{n-k}$.

4′. *Cautious Dynamic Pareto optimality (CDPO)*. A dynamic allocation mechanism satisfies CDPO if at every step k, the allocation $\boldsymbol{A}^k$ returned by the mechanism is not Pareto dominated by any other allocation $\boldsymbol{A}'$ over the same k agents that is EF-extensible.

In other words, a mechanism satisfies CDPO if at every step it selects an allocation that is at least as generous as any allocation that can ultimately guarantee EF, irrespective of future demands.

At first glance, it may not be obvious that CDPO is indeed a relaxation of DPO (i.e., that CDPO is implied by DPO). However, note that DPO requires a mechanism to allocate at least a k/n fraction of at least one resource r^* in the allocation $\boldsymbol{A}^k$ at any step k, and thus to allocate at least a $1/n$ fraction of that resource to some agent i. Any alternative allocation that Pareto dominates $\boldsymbol{A}^k$ must also allocate at least a $1/n$ fraction of r^* to agent i. Consequently, in order to ensure an EF extension over all n agents when all the future demands are identical to the demand of agent i, the alternative allocation must allocate at most a k/n fraction of r^*, as each future agent may also require at least a $1/n$ fraction of r^* to avoid envying agent i. It follows that the alternative allocation cannot Pareto dominate $\boldsymbol{A}^k$. Thus, the mechanism satisfies CDPO.

Recall that DYNAMIC DRF extends the water-filling idea of the static DRF mechanism [99] to our dynamic setting. DYNAMIC DRF is unable to satisfy the original EF, because — to satisfy DPO — at every step k it needs to allocate resources until a k/n fraction of some resource is allocated. We wish to modify DYNAMIC DRF to focus only on competing with EF-extensible allocations, in a way that achieves CDPO and EF (as well as other properties).

The main technical challenge is checking when an allocation at step k violates EF-extensibility. Indeed, there are uncountably many possibilities for the future demands $\boldsymbol{d}_{>k}$ over which an EF extension needs to be guaranteed by an EF-extensible allocation! Of course, checking all the possibilities explicitly is not feasible. Ideally, we would like to check only a small number of possibilities. The following lemma establishes that it is sufficient to verify that an EF extension exists under the assumption that all future agents will have the same demand vector that is moreover identical to the demand vector of one of the present agents.

Lemma 5.5. *Let k be the number of present agents, $\boldsymbol{d}_{\leqslant k}$ be the demands reported by the present agents, and $\boldsymbol{A}$ be an EF allocation over the k present agents. Then $\boldsymbol{A}$ is EF-extensible if and only if there exists an EF extension of $\boldsymbol{A}$ over all n agents with demands $\boldsymbol{d} = (\boldsymbol{d}_{\leqslant k}, \boldsymbol{d}_{>k})$ for all future demands $\boldsymbol{d}_{>k} \in \mathcal{D}'$, where $\mathcal{D}' = \{\langle \boldsymbol{d}_1 \rangle^{n-k}, \langle \boldsymbol{d}_2 \rangle^{n-k}, \ldots, \langle \boldsymbol{d}_k \rangle^{n-k}\}$.*

To prove this lemma, we first introduce the notion of the *minimum EF extension*. Intuitively, the minimum EF extension is the "smallest" EF extension (allocating the least resources) of a given EF allocation to a larger set of agents. Formally, let $\boldsymbol{A}$ be an EF allocation over a set of agents $S \subseteq N$ and $\boldsymbol{A}^*$ be an EF extension of A to a set of agents $T \subseteq N$ ($S \subseteq T$). Then $\boldsymbol{A}^*$ is called the minimum EF extension of $\boldsymbol{A}$ to T if for any EF extension $\boldsymbol{A}'$ of $\boldsymbol{A}$ to T, we have that $\boldsymbol{A}'$ is an extension of $\boldsymbol{A}^*$. We show that the minimum EF extension exists and exhibits a simple structure.

Lemma 5.6. *Let $\boldsymbol{A}$ be an EF allocation over a set of agents $S \subseteq N$ and let x_i be the dominant share of agent $i \in S$ in $\boldsymbol{A}$. Let T be such that $S \subseteq T \subseteq N$ and let $\boldsymbol{A}^*$ be an allocation over T with x_i^* as the dominant share of agent $i \in T$. Let $x_i^* = x_i$ for all $i \in S$, and $x_i^* = \max_{j \in S} y_i^j$ for all $i \in T \setminus S$, where $y_i^j = x_j \cdot \min_{r \in R} d_{jr}/d_{ir}$. Then $\boldsymbol{A}^*$ is a minimum EF extension of $\boldsymbol{A}$ to T.*

Proof. For agent i with dominant share x_i to avoid envying agent j with dominant share x_j, there must exist $r \in R$ such that $x_i \cdot d_{ir} \geqslant x_j \cdot d_{jr}$, that is, $x_i \geqslant x_j \cdot d_{jr}/d_{ir}$. It follows that $x_i \geqslant x_j \cdot \min_{r \in R} d_{jr}/d_{ir}$, and thus the minimum dominant share is given by $y_i^j = x_j \cdot \min_{r \in R} d_{jr}/d_{ir}$. Now it is easy to argue that any EF extension $\boldsymbol{A}'$ of $\boldsymbol{A}$ over T must allocate at least an x_i^* dominant share to any agent $i \in T$, for both $i \in S$ and $i \in T \setminus S$, and thus $\boldsymbol{A}'$ must be an extension of $\boldsymbol{A}^*$.

It remains to prove that $\boldsymbol{A}^*$ is EF. First we prove an intuitive result regarding the minimum dominant share agent i needs to avoid envying agent j, namely y_i^j. We claim that for every $r \in R$,

$$y_i^j \cdot d_{ir} \leqslant x_j \cdot d_{jr}. \tag{5.2}$$

Formally, for any $r \in R$,

$$y_i^j \cdot d_{ir} = x_j \cdot \min_{r' \in R} \frac{d_{jr'}}{d_{ir'}} \cdot d_{ir} \leqslant x_j \cdot \frac{d_{jr}}{d_{ir}} \cdot d_{ir} = x_j \cdot d_{jr}.$$

Therefore, to prevent agent i from envying agent j, we need to allocate at least an $x_j \cdot d_{jr}$ fraction of resource r to agent i for some $r \in R$. Next we show that $\boldsymbol{A}^*$ is EF, i.e., no agent i envies any agent j in $\boldsymbol{A}^*$. We consider four cases.

Case 1: $i \in S$ and $j \in S$. This case is trivial as $\boldsymbol{A}^*$ is identical to $\boldsymbol{A}$ over S and $\boldsymbol{A}$ is EF.

Case 2: $i \in T \setminus S$ and $j \in S$. This case is also trivial because i receives at least a y_i^j fraction of its dominant resource.

Case 3: $i \in S$ and $j \in T \setminus S$. We must have $x_j = y_j^t$ for some $t \in S$. Agent i does not envy agent t in $\boldsymbol{A}$, and hence in $\boldsymbol{A}^*$. Thus, there exists a resource $r \in R$ such that $A_{ir}^* \geqslant A_{tr}^* \geqslant A_{jr}^*$, where the last step follows from Equation (5.2). Thus, agent i does not envy agent j.

Case 4: $i \in T \setminus S$ and $j \in T \setminus S$. Similarly to Case 3, let $x_j = y_j^t$ for some $t \in S$. Now $x_i \geqslant y_i^t$, so agent i does not envy agent t in $\boldsymbol{A}^*$. Thus, there exists a resource r such that $A_{ir}^* \geqslant A_{tr}^* \geqslant A_{jr}^*$, where again the last step follows from Equation (5.2).

Therefore, $\boldsymbol{A}^*$ is an EF extension of $\boldsymbol{A}$ over T and we have already established that any EF extension of $\boldsymbol{A}$ over T must be an extension of $\boldsymbol{A}^*$. We conclude that $\boldsymbol{A}^*$ is a minimum EF extension of $\boldsymbol{A}$ over T. ■ (Proof of Lemma 5.6)

It is not hard to see from the construction of the minimum EF extension that it not only exists, it is unique. We are now ready to prove Lemma 5.5.

Proof of Lemma 5.5. The "only if" direction of the proof is trivial. To prove the "if" part, we prove its contrapositive. Assume that there exist future demand vectors $\widehat{\boldsymbol{d}}_{>k} \in \mathcal{D}^{n-k}$ such that there does not exist any EF extension of $\boldsymbol{A}$ to N with demands $\widehat{\boldsymbol{d}} = (\boldsymbol{d}_{\leqslant k}, \widehat{\boldsymbol{d}}_{>k})$. We want to show that there exists $\boldsymbol{d}'_{>k} \in \mathcal{D}'$ for which there is no EF extension as well.

Let $K = \{1, \ldots, k\}$ and $N \setminus K = \{k+1, \ldots, n\}$. Denote the minimum EF extension of $\boldsymbol{A}$ to N with demands $\widehat{\boldsymbol{d}}$ by $\boldsymbol{A}^*$. Let the dominant share of agent $i \in K$ in $\boldsymbol{A}$ be x_i and the dominant share of agent $j \in N$ in $\boldsymbol{A}^*$ be x_j^*.

No EF extension of $\boldsymbol{A}$ over N with demands $\widehat{\boldsymbol{d}}$ is feasible, hence $\boldsymbol{A}^*$ must be infeasible too. Therefore, there exists a resource r such that $\sum_{i=1}^n x_i^* \cdot d_{ir} > 1$. Note that for every agent $j \in N \setminus K$, there exists an agent $i \in K$ such that $x_j^* = x_i \cdot \min_{r' \in R} d_{ir'}/d_{jr'}$, and hence $x_j^* \cdot d_{jr} \leqslant x_i \cdot d_{ir}$ by Equation (5.2). Taking the maximum over $i \in K$, we get that $x_j^* \cdot d_{jr} \leqslant \max_{i \in K} (x_i \cdot d_{ir})$ for every agent $j \in N \setminus K$. Taking $t \in \arg\max_{i \in K} (x_i \cdot d_{ir})$,

$$\begin{aligned} 1 < \sum_{i=1}^{n} x_i^* \cdot d_{ir} &= \sum_{i=1}^{k} x_i^* \cdot d_{ir} + \sum_{i=k+1}^{n} x_i^* \cdot d_{ir} \\ &\leqslant \sum_{i=1}^{k} x_i \cdot d_{ir} + (n-k) \cdot x_t \cdot d_{tr}. \end{aligned}$$

Consider the case where $\boldsymbol{d}'_{>k} = \langle \boldsymbol{d}_t \rangle^{n-k} \in \mathcal{D}'$. The minimum EF extension $\boldsymbol{A}'$ of $\boldsymbol{A}$ to N with demands $\boldsymbol{d}' = \left\langle \boldsymbol{d}_{\leqslant k}, \boldsymbol{d}'_{>k} \right\rangle$ allocates an x_i dominant share to every $i \in K$ (same as $\boldsymbol{A}$) and allocates exactly an x_t dominant share to every $j \in N \setminus K$. Thus, the fraction of resource r allocated in $\boldsymbol{A}'$ is $\sum_{i=1}^k x_i \cdot d_{ir} + (n-k) \cdot x_t \cdot d_{tr} > 1$, implying that the minimum EF extension of $\boldsymbol{d}'_{>k}$ is infeasible. We conclude that there is no feasible EF extension for $\boldsymbol{d}'_{>k}$, as required. ■ (Proof of Lemma 5.5)

The equivalent condition of Lemma 5.5 provides us with $k \cdot m$ linear constraints that can be checked to determine whether an allocation over k agents is EF-extensible. Using this machinery, we can write down a "small" linear program (LP) that begins with the allocation chosen in the previous step (recall that the allocations are irrevocable), gives agent k a jump start so that it does not envy agents 1 through $k-1$, and then uses water-filling to allocate resources similarly to DYNAMIC DRF, but subject to the constraint that the allocation stays EF-extensible. This intuition is formalized via the mechanism CAUTIOUS LP, which is given as Algorithm 7.

ALGORITHM 7: CAUTIOUS LP

Data: Demands $\boldsymbol{d}$
Result: Allocation $\boldsymbol{A}^k$ at each step k
$k \leftarrow 1$;
while $k \leqslant n$ **do**
  $\{x_i^k\}_{i=1}^k \leftarrow$ Solution of the LP in the box below;
  $A_{ir}^k \leftarrow x_i^k \cdot d_{ir}, \forall i \leqslant k$;
  $k \leftarrow k+1$
end

Maximize M^k
subject to
$x_i^k \geqslant M^k, \forall i \leqslant k$
$x_i^k \geqslant x_i^{k-1}, \forall i \leqslant k-1$
$x_k^k \geqslant \max_{i \leqslant k-1} \left(x_i^{k-1} \cdot \min_{r \in R} d_{ir}/d_{kr} \right)$
$\sum_{i=1}^k x_i^k \cdot d_{ir} + (n-k) \cdot x_t^k \cdot d_{tr} \leqslant 1, \forall t \leqslant k, r \in R$

The mechanism's third LP constraint jump-starts agent k to a level where it does not envy earlier agents, and the fourth LP constraint is derived from Lemma 5.5. To see why the mechanism satisfies CDPO, observe that if at any step k there is an EF-extensible allocation $\boldsymbol{A}'$ that Pareto dominates the allocation $\boldsymbol{A}^k$ returned by the mechanism, then (by Lemma 5.5) $\boldsymbol{A}'$ must also satisfy the LP at step k. However, it can be shown that no allocation from the feasible region of the LP can Pareto dominate $\boldsymbol{A}^k$. Indeed, if an allocation from the feasible region did dominate $\boldsymbol{A}^k$, we could redistribute some of the resources of the agent that is strictly better off to obtain a feasible allocation with a value of M^k that is higher than the optimal solution. It is also easy to see why intuitively CAUTIOUS LP is EF: the initial allocation to agent k achieves an EF allocation over the k agents, and water-filling preserves EF because it always allocates to agents with minimum dominant share. It is equally straightforward to show that CAUTIOUS LP also satisfies SI. Establishing SP requires some work, but the proof is mainly a modification of the proof of Theorem 5.2. We are therefore able to establish the following theorem, which formalizes the guarantees given by CAUTIOUS LP.

Theorem 5.3. CAUTIOUS LP *satisfies SI, EF, CDPO, and SP, and can be implemented in polynomial time.*

Proof. The proof is along the lines of the proof of Theorem 5.2. For now, assume that the LP is feasible at each step and thus the mechanism does return an allocation at each step (we show this below). In the LP at step k, let

$$E^k = \max_{i \leqslant k-1} \left(x_i^{k-1} \cdot \min_{r \in R} d_{ir} / d_{kr} \right).$$

Intuitively, E^k is the jump start that agent k requires at the beginning of step k to be envy free of the allocations of agents 1 through $k-1$ from step $k-1$.

Proof of CDPO: First we show that CAUTIOUS LP satisfies CDPO. Assume for contradiction, that at some step $k \in \{1, \dots, n\}$, an alternative EF-extensible allocation $\boldsymbol{A}'$ over the k present agents Pareto dominates the allocation $\boldsymbol{A}^k$ returned by the mechanism. Let x'_i be the dominant share of agent i in $\boldsymbol{A}'$, for $i \leqslant k$. Since $\boldsymbol{A}'$ Pareto dominates $\boldsymbol{A}^k$, we have that $x'_i \geqslant x_i^k$ for every $i \leqslant k$. This trivially implies that $\boldsymbol{A}'$ also satisfies the first three constraints of the LP at step k. Moreover, since $\boldsymbol{A}'$ is EF-extensible, it also satisfies the fourth constraint of the LP at step k as the fourth constraint only requires EF extension to exist in specific cases (in particular, it requires the minimum EF extension and thus any EF extension over all n agents to exist when all future demand vectors are identical to the demand vector of some present agent). Thus, $\boldsymbol{A}'$ is in the feasible region of the LP and Pareto dominates an optimal solution $\boldsymbol{A}^k$. Now, taking back the extra resources that $\boldsymbol{A}'$ allocates to agents compared to $\boldsymbol{A}^k$ shows that the fourth constraint is not tight in $\boldsymbol{A}^k$ for any value of t and r (the assumption of strictly positive demands is crucial here). However, this implies that in the allocation $\boldsymbol{A}^k$, every x_i^k and correspondingly M^k can be increased by a sufficiently small quantity while still satisfying the LP at step k, which contradicts the optimality of $\boldsymbol{A}^k$. Thus, no alternative EF-extensible allocation can Pareto dominate the allocation given by the mechanism at any step, i.e., CAUTIOUS LP satisfies CDPO.

Proof of SI: Next, we show that CAUTIOUS LP satisfies SI. We show this by induction over step k. For the base case $k = 1$, it is easy to show that setting $x_1^1 = 1/n$ and $M^k = 1/n$ satisfies the LP at step 1; it trivially satisfies the first three constraints of the LP and for the fourth constraint, observe that

$$\frac{1}{n} \cdot d_{ir} + (n-1) \cdot \frac{1}{n} \cdot d_{ir} = d_{ir} \leqslant 1, \forall r \in R.$$

Therefore, in the optimal solution, $M^1 \geqslant 1/n$ and thus $x_1^1 \geqslant 1/n$ (in fact, equality holds).

Now consider any step $k \in \{2, \dots, n\}$. As our induction hypothesis, we assume that $x_i^t \geqslant 1/n$ for all agents $i \leqslant t$, at every step $t \leqslant k-1$. We want to show that $x_i^k \geqslant 1/n$ for all agents $i \leqslant k$. Consider two cases.

1. $E^k \geqslant 1/n$. Observe that $x_i^{k-1} \geqslant 1/n$ for all $i \leqslant k-1$ due to the induction hypothesis. Thus, using the second and the third constraints of the LP at step k, we have $x_i^k \geqslant 1/n$ for all $i \leqslant k$.
2. $E^k < 1/n$. We first show that $x_i^k = x_i^{k-1}$ for $i \leqslant k-1$, $x_k^k = 1/n$ and $M^k = 1/n$ is in the feasible region of the LP at step k. Note that this assignment trivially satisfies the first three constraints of the LP.

For the fourth constraint, fix any $r \in R$. Define $T_r = \max_{i \leqslant k-1} x_i^{k-1} \cdot d_{ir}$. First, we show that $\sum_{i=1}^{k-1} x_i^{k-1} \cdot d_{ir} \leqslant 1 - (n-k+1) \cdot \max(T_r, 1/n)$. To see this, note that $\{x_i^{k-1}\}_{i=1}^{k-1}$ satisfies the LP at step $k-1$ and, in particular, the fourth constraint of the LP. Therefore,

$$\sum_{i=1}^{k-1} x_i^{k-1} \cdot d_{ir} + (n-k+1) \cdot T_r \leqslant 1 \Longrightarrow \sum_{i=1}^{k-1} x_i^{k-1} \cdot d_{ir} \leqslant 1 - (n-k+1) \cdot T_r.$$

Now we prove that

$$\sum_{i=1}^{k-1} x_i^{k-1} \cdot d_{ir} \leqslant 1 - (n-k+1) \cdot 1/n = (k-1)/n. \tag{5.3}$$

Suppose for contradiction that the left hand side is more than $(k-1)/n$. Then, by the pigeonhole principle, there exists some agent $i \leqslant k-1$ such that $x_i^{k-1} \cdot d_{ir} \geqslant 1/n$, and thus $T_r \geqslant 1/n$. But we have already shown that

$$\sum_{i=1}^{k-1} x_i^{k-1} \cdot d_{ir} \leqslant 1 - (n-k+1) \cdot T_r \leqslant 1 - (n-k+1) \cdot 1/n = (k-1)/n,$$

contradicting our assumption; this establishes (5.3). Thus, we have that

$$\sum_{i=1}^{k-1} x_i^{k-1} \cdot d_{ir} \leqslant 1 - (n-k+1) \cdot \max\left(T_r, \frac{1}{n}\right).$$

Finally, we show that in the fourth constraint of the LP, $x_t^k \cdot d_{tr} \leqslant \max(T_r, 1/n)$. To see this, observe that for $t \leqslant k-1$, $x_t^k \cdot d_{tr} = x_t^{k-1} \cdot d_{tr} \leqslant T_r$ and for $t = k$, $x_t^k \cdot d_{tr} = 1/n \cdot d_{kr} \leqslant 1/n$. Thus, the fourth constraint of the LP is satisfied for every $t \leqslant k$ and every $r \in R$.

We have established that CAUTIOUS LP satisfies SI. Our next goal is to prove that the mechanism also satisfies EF and SP. As in the proof of Theorem 5.2, we first establish several useful lemmas about the allocations returned by CAUTIOUS LP. In the proof below, M^k and x_i^k refer to the *optimal* solution of the LP in step k.

We begin with the following lemma (similar to Lemma 5.1), which essentially shows that if an agent is allocated some resources in step k using water-filling (in addition to the jump-start to E^k for agent k), then the agent's dominant share after the step would be the minimum among the present agents.

Lemma 5.7. *At every step $k \in \{1, \ldots, n\}$, it holds that $x_i^k = \max(M^k, x_i^{k-1})$ for all agents $i \leqslant k-1$, and $x_k^k = \max(M^k, E^k)$.*

Proof. Consider any step $k \in \{1, \ldots, n\}$. From the first three constraints of the LP, it is evident that $x_i^k \geqslant M^k$ for all $i \leqslant k$, $x_i^k \geqslant x_i^{k-1}$ for all $i \leqslant k-1$ and $x_k^k \geqslant E^k$. Thus, $x_i^k \geqslant \max(M^k, x_i^{k-1})$ for all $i \leqslant k-1$ and $x_k^k \geqslant \max(M^k, E^k)$.

Suppose for contradiction that a strict inequality holds for some agent $i \leqslant k$. Then x_i^k can be reduced by a sufficiently small $\varepsilon > 0$ without violating any constraints. This makes the third constraint of the LP loose by at least $\varepsilon \cdot d_{ir}$, for every resource $r \in R$. Consequently, the values of x_j^k for $j \neq i$ and M^k can be increased by a sufficiently small $\delta > 0$ without violating the third constraint of the LP. Finally, ε (and correspondingly δ) can be chosen to be small enough so that $x_i^k \geqslant M^k$ is not violated. It follows that the value of M^k can be increased, contradicting the optimality of M^k. ■ (Proof of Lemma 5.7)

Next, we formulate the equivalent of Lemma 5.2 as two separate lemmas. First we show that if an agent has greater or equal dominant share than another agent in some step (where both are present), then the order is preserved in future steps. Next we show that at each step k, the dominant shares of agents 1 through k are monotonically non-increasing with their time of arrival, except for agents that have not received any resources apart from their jump-start.

Lemma 5.8. *For any agents $i, j \in N$ and any step $k \geqslant \max(i, j)$ (i.e., both agents are present at step k), $x_i^k \geqslant x_j^k$ implies that $x_i^t \geqslant x_j^t$ for all $t \geqslant k$.*

Proof. Fix any two agents $i, j \in N$ and step $k \geqslant \max(i, j)$ such that $x_i^k \geqslant x_j^k$. We use induction on t. The result trivially holds for $t = k$. Consider any $t > k$ and assume the result holds for step $t - 1$. Then, since $t > k \geqslant \max(i, j)$ we know that $x_i^t = \max(x_i^{t-1}, M^t) \geqslant \max(x_j^{t-1}, M^t) = x_j^t$, where the first and the last transitions follow from Lemma 5.7 and the second transition follows from our induction hypothesis. ■ (Proof of Lemma 5.8)

Lemma 5.9. *For all agents $i, j \in N$ such that $i < j$ and any step $k \geqslant j$, we have that either i) $x_i^k \geqslant x_j^k$ or ii) $x_j^k = x_j^j = E^j$.*

Proof. Fix any two agents $i, j \in N$ such that $i < j$ and any step $k \geqslant j$. Note that $x_j^k \geqslant x_j^j \geqslant E^j$, where the first inequality is due to irrevocability of resources and the last inequality is due to Lemma 5.7. If $x_j^k = E^j$, then the lemma trivially holds. Assume $x_j^k > E^j$. Consider the first step t where $x_j^t > E^j$ (thus $j \leqslant t \leqslant k$). If $t = j$, then we have $x_j^j > E^j$. If $t > j$, then we have $x_j^t > x_j^{t-1}$ since $x_j^{t-1} = E^j$ by the definition of t. In any case, Lemma 5.7 implies that $x_j^t = M^t \leqslant x_i^t$. Thus we have $x_j^t \leqslant x_i^t$ and now Lemma 5.8 implies that $x_j^k \leqslant x_i^k$. ■ (Proof of Lemma 5.9)

We now consider the equivalent of Lemma 5.3 from the proof of Theorem 5.2, and observe that there are two cases. If agent j has greater dominant share than agent i at some step, then either j arrived before i and j has not been allocated any resources since i arrived (as we had previously), or j arrived after i and has not been allocated any resources apart from its jump-start.

Lemma 5.10. *For any agents $i, j \in N$ and any step $k \geqslant \max(i, j)$ (i.e., both agents are present), $x_j^k > x_i^k$ implies that either i) $j < i$ and $x_j^k = x_j^{i-1}$, or ii) $j > i$ and $x_j^k = x_j^j = E^j$.*

Proof. Fix any two agents $i, j \in N$ and any step $k \geqslant \max(i, j)$ such that $x_j^k > x_i^k$. Note that if $j > i$ then Lemma 5.9 implies that $x_j^k = x_j^j = E^j$ and the result holds trivially. Now assume $j < i$.

Suppose for contradiction that $x_j^k > x_j^{i-1}$ (it cannot be smaller because allocations are irrevocable). Then there exists a step $t \in \{i, \ldots, k\}$ such that $x_j^t > x_j^{t-1}$. Therefore, Lemma 5.7 implies that $x_j^t = M^t \leqslant x_i^t$. Using Lemma 5.8 this shows that $x_j^k \leqslant x_i^k$, which is a contradiction to the assumption that $x_j^k > x_i^k$. Thus we have $x_j^k = x_j^{i-1}$, as required. ■ (Proof of Lemma 5.10)

Finally, we establish an additional lemma which will be helpful in proving SP. For agents i, j such that $j > i$, if the jump-start E^j for agent j requires allocating agent j greater dominant share than agent i had in step $j - 1$, then clearly the jump-start must have been due to agent j envying some agent $l \neq i$, and l must have greater dominant share than i in step $j - 1$. But then using Lemma 5.9 and extending the argument, we can eventually trace this back to an agent $t < i$. We show that we can find such $t < i$ such that the jump-start of the original agent j was actually due to agent j envying agent t.

Lemma 5.11. *For any agents $i, j \in N$ such that $j > i$, $E^j > x_i^j$ implies that $E^j = x_t^{j-1} \cdot \min_{r \in R} d_{tr}/d_{jr}$, for some agent $t < i$.*

Proof. Fix any agent $i \in N$. We use induction over $j \in \{i+1, \ldots, n\}$. First, we prove several implications that hold for any agent $j > i$. Recall that $E^j = \max_{p<j}(x_p^{j-1} \cdot \min_{r \in R} d_{pr}/d_{jr})$. Thus, we have $E^j = x_l^{j-1} \cdot \min_{r \in R} d_{lr}/d_{jr}$ for some agent $l < j$. But it does not follow from the definition that we can take $l < i$. Observe that

$$x_l^{j-1} \geqslant x_l^{j-1} \cdot \min_{r \in R} d_{tr}/d_{jr} = E^j > x_i^j \geqslant x_i^{j-1}, \tag{5.4}$$

where the first transition holds since $\min_{r \in R} d_{tr}/d_{jr}$ is at most 1 (consider the dominant resource of agent j), the third transition is the assumption of the lemma and the last transition holds since allocations are irrevocable.

Now we have three cases. If $l < i$, then we are done. Further, $l \neq i$ since Equation (5.4) shows that $x_l^{j-1} > x_i^{j-1}$. Now assume that $l > i$. Note that this case cannot appear in the base case $j = i + 1$ since $l < j$. Therefore, the argument given above already shows that the lemma holds for the base case $j = i + 1$. By our induction hypothesis, we assume that the lemma holds for agent $l < j$. Now since $l > i$ and $x_l^{j-1} > x_i^{j-1}$, Lemma 5.10 implies that $x_l^{j-1} = x_l^l = E^l$ and thus $E^l > x_i^{j-1} \geqslant x_i^l$ where $x_i^{j-1} \geqslant x_i^l$ because $l < j$ and allocations are irrevocable. Due to our induction hypothesis, there exists $t < i$ such that $E^l = x_t^{l-1} \cdot \min_{r \in R} d_{tr}/d_{lr}$. We prove that $E^j = x_t^{j-1} \cdot \min_{r \in R} d_{tr}/d_{jr}$. Indeed,

$$E^j = x_l^{j-1} \cdot \min_{r \in R} d_{lr}/d_{jr}$$

$$
\begin{aligned}
&= E^l \cdot \min_{r\in R} d_{lr}/d_{jr} \\
&= x_t^{l-1} \cdot \min_{r\in R} d_{tr}/d_{lr} \cdot \min_{r\in R} d_{lr}/d_{jr} \\
&\leqslant x_t^{l-1} \cdot \min_{r\in R} d_{tr}/d_{jr} \\
&\leqslant x_t^{j-1} \cdot \min_{r\in R} d_{tr}/d_{jr} \leqslant E^j.
\end{aligned}
$$

Here, the fourth transition is true because for any $r' \in R$,

$$\frac{d_{tr'}}{d_{jr'}} = \frac{d_{tr'}}{d_{lr'}} \cdot \frac{d_{lr'}}{d_{jr'}} \geqslant \min_{r\in R} \frac{d_{tr}}{d_{lr}} \cdot \min_{r\in R} \frac{d_{lr}}{d_{jr}}.$$

Taking minimum over all $r' \in R$, we get that $\min_{r\in R} d_{tr}/d_{jr} \geqslant \min_{r\in R} d_{tr}/d_{lr} \cdot \min_{r\in R} d_{lr}/d_{jr}$. The last transition holds due to the definition of E^j. Now it is trivial to see that we must have equality at every step, so $E^j = x_t^{j-1} \cdot \min_{r\in R} d_{tr}/d_{jr}$ for $t < i$, as required. ■ (Proof of Lemma 5.11)

Proof of LP Feasibility and EF: Now we use an inductive argument to simultaneously show that the LP of CAUTIOUS LP is feasible at every step and that CAUTIOUS LP satisfies EF. Consider the following induction hypothesis: the LP at step t is feasible and the allocation $\boldsymbol{A}^t$ returned by the mechanism at step t is EF. For the base case $t = 1$, the LP is trivially feasible and the allocation $\boldsymbol{A}^1$ is also trivially EF. Assume that the hypothesis holds for $t = k - 1$ for some step $k \in \{2, \ldots, n\}$. We want to show that the hypothesis holds for step k.

For feasibility, we show that the allocation $\boldsymbol{A}^*$ given by $x_i^k = x_i^{k-1}$ for $i \leqslant k - 1$ and $x_k^k = E^k$ along with $M^k = 0$ satisfies the LP at step k. Clearly, it satisfies the first three constraints of the LP. To see why it satisfies the fourth constraint, note that $\boldsymbol{A}^{k-1}$ is an EF allocation due to our induction hypothesis. Moreover, it satisfies the LP at step $k - 1$, in particular, the fourth constraint of the LP. Hence Lemma 5.5 implies that $\boldsymbol{A}^{k-1}$ must be an EF-extensible allocation. Let $\boldsymbol{d}_k$ denote the demand reported by agent k in step k and let $\boldsymbol{d}_{>k} \in \mathcal{D}^{n-k}$. Then any EF extension of $\boldsymbol{A}^{k-1}$ over all n agents with future demands $(\boldsymbol{d}_k, \boldsymbol{d}_{>k})$ is an EF extension of $\boldsymbol{A}^*$ over all n agents with future demands $\boldsymbol{d}_{>k}$. Since this holds for any $\boldsymbol{d}_{>k} \in \mathcal{D}^{n-k}$, $\boldsymbol{A}^*$ is EF-extensible and hence satisfies the fourth constraint of the LP. We conclude that the LP is feasible at step k.

Now we want to show that the allocation $\boldsymbol{A}^k$ is an EF allocation. Intuitively, we can see that the mechanism starts from $\boldsymbol{A}^*$ which is EF (it is a minimum EF extension), and then uses water-filling to allocate more resources in a way that preserves EF. Formally, note that the dominant shares allocated to agents in $\boldsymbol{A}^k$ are given by Lemma 5.7. Take any two agents $i, j \leqslant k$. We want to show that agent i does not envy agent j in step k. Denote the dominant share of an agent l in $\boldsymbol{A}^*$ by x_l^*, i.e., $x_l^* = x_l^{k-1}$ for $l \leqslant k - 1$ and $x_k^* = E^k$. It holds that

$$x_i^k = \max\left(x_i^*, M^k\right) \geqslant \max\left(x_j^* \cdot \min_{r\in R} \frac{d_{jr}}{d_{ir}}, M^k\right) \geqslant \max\left(x_j^*, M^k\right) \cdot \min_{r\in R} \frac{d_{jr}}{d_{ir}}$$

$$= x_j^k \cdot \min_{r \in R} d_{jr}/d_{ir},$$

where the first and the last transitions follow from Lemma 5.7, the second transition holds since the allocation $\boldsymbol{A}^*$ is EF, and the third transition holds since the quantity $\min_{r \in R} d_{jr}/d_{ir}$ is at most 1. Thus, $\boldsymbol{A}^k$ is EF. By induction, it holds that the LP of CAUTIOUS LP is feasible at every step and CAUTIOUS LP is EF.

Proof of SP: Our last task is to prove that CAUTIOUS LP is SP. Suppose for contradiction that an agent $i \in N$ can report an untruthful demand vector $\boldsymbol{d}_i'$ such that the agent is strictly better off in at least one step. Let k be the first such step. Denote by $\widehat{x}_j^t$ the dominant share of an agent j at step t with manipulation (for agent i, this is the share of the dominant resource of the untruthful demand vector) and, similarly, denote by $\widehat{M}^t$ the value of M^t in the optimal solution of the LP of step t with manipulation.

Lemma 5.12. *$\widehat{x}_j^k \geqslant x_j^k$ for every agent $j \leqslant k$.*

Proof. Fix any agent $j \leqslant k$. We provide a case by case analysis and show that the lemma holds in each case.

1. $x_j^k \leqslant x_i^k$. In this case, we have

$$x_j^k \leqslant x_i^k < \widehat{x}_i^k = \widehat{M}^k \leqslant \widehat{x}_j^k.$$

   The second transition holds because if $\widehat{x}_i^k \leqslant x_i^k$ then agent i could not be better off as the share of the dominant resource of its true demand vector that it receives with manipulation would be no more than it received without manipulation. To justify the third transition, note that agent i must be allocated some resources at step k with manipulation. If $k = i$, then note that since E^i only depends on the allocation at step $i-1$ which is not affected due to manipulation by agent i, we have $\widehat{E}^i = E^i \leqslant x_i^i < \widehat{x}_i^i$ and Lemma 5.7 implies that $\widehat{x}_i^i = \widehat{M}^i$. If $k > i$ and $\widehat{x}_i^k \neq \widehat{M}^k$, then Lemma 5.7 implies that $\widehat{x}_i^k = \widehat{x}_i^{k-1}$, but then $u_i(\widehat{\boldsymbol{A}}_i^{k-1}) = u_i(\widehat{\boldsymbol{A}}_i^k) > u_i(\boldsymbol{A}_i^k) \geqslant u_i(\boldsymbol{A}_i^{k-1})$, where $\widehat{\boldsymbol{A}}_i^k$ is the allocation to agent i at step k with manipulation. That is, agent i would have been better off with manipulation in step $k-1$, which is a contradiction since k is the first such step. The last transition holds because $\widehat{x}_j^k$ satisfies the first constraint of the LP of step k with manipulation.
2. $x_j^k > x_i^k$. For this, we have three sub-cases.
   (a) $j < i$. Then we have $x_j^k = x_j^{i-1} = \widehat{x}_j^{i-1} \leqslant \widehat{x}_j^k$, where the first transition follows due to Lemma 5.10, the second transition holds because manipulation by agent i does not affect the allocations at step $i-1$, and the third transition follows since allocations are irrevocable.
   (b) $j = i$. This cannot happen since we have assumed $x_j^k > x_i^k$ in this case.
   (c) $j > i$. Since $x_j^k > x_i^k$, Lemma 5.10 implies that $x_j^k = x_j^j = E^j$, so $E^j > x_i^k$. Now using Lemma 5.11, $E^j = x_t^{j-1} \cdot \min_{r \in R} d_{tr}/d_{jr}$ for some $t < i$. Then, $x_t^{j-1} \geqslant x_t^{j-1} \cdot \min_{r \in R} d_{tr}/d_{jr} = E^j > x_i^k \geqslant x_i^{j-1}$, where the first transition follows since

$\min_{r \in R} d_{tr}/d_{jr}$ is at most 1 and the last transition follows since allocations are irrevocable. Now Lemma 5.10 implies that $x_t^{j-1} = x_t^{i-1}$. Putting all the pieces together,

$$\begin{aligned} x_j^k = E^j = x_t^{j-1} \cdot \min_{r \in R} d_{tr}/d_{jr} = x_t^{i-1} \cdot \min_{r \in R} d_{tr}/d_{jr} = x_t^{j-1} \cdot \min_{r \in R} d_{tr}/d_{jr} \\ \leqslant \widehat{x}_t^{j-1} \cdot \min_{r \in R} d_{tr}/d_{jr} \leqslant \widehat{E}^j \leqslant \widehat{x}_j^j \leqslant \widehat{x}_j^k, \end{aligned}$$

where the fifth transition follows since manipulation by agent i does not change the allocation at step $i-1$, the sixth transition follows due to the definition of $\widehat{E}^j$ (which is the value of E^j after manipulation), the seventh transition follows due to the third constraint of the LP at step j after manipulation, and the last transition follows since allocations are irrevocable.

Thus, we conclude that $\widehat{x}_j^k \geqslant x_j^k$ for all agents $j \leqslant k$. ■ (Proof of Lemma 5.12)

Now, in the optimal solution of the LP at step k without manipulation (i.e., in $\boldsymbol{A}^k$), the fourth constraint must be tight for some $t \leqslant k$ and $r \in R$ (otherwise x_j^k for every $j \leqslant k$ and M^k can be increased, contradicting the optimality of M^k). Thus,

$$\sum_{j=1}^{k} x_j^k \cdot d_{jr} + (n-k) \cdot x_t^k \cdot d_{tr} = 1.$$

Now consider the fourth constraint of the LP at step k after manipulation for the same values of t and r. For simplicity of notation, let $d'_{jr} = d_{jr}$ for $j \neq i$. Then,

$$\sum_{j=1}^{k} \widehat{x}_j^k \cdot d'_{jr} + (n-k) \cdot \widehat{x}_t^k \cdot d'_{tr} > \sum_{j=1}^{k} x_j^k \cdot d_{jr} + (n-k) \cdot x_t^k \cdot d_{tr} = 1.$$

To justify the inequality, note that $\widehat{x}_i^k \cdot d'_{ir} > x_i^k \cdot d_{ir}$ by Equation (5.1) (as agent i is strictly better off), and for any $j \leqslant k$ such that $j \neq i$, $\widehat{x}_j^k \cdot d'_{jr} = \widehat{x}_j^k \cdot d_{jr} \geqslant x_j^k \cdot d_{jr}$ by Lemma 5.12. However, this shows that the allocation at step k with manipulation violates the fourth constraint of the LP, which is impossible. Hence, a successful manipulation is impossible, that is, CAUTIOUS LP is SP.

Finally, note that at every step the LP has $O(n)$ variables and $O(n \cdot m)$ constraints, and there are n such steps. Hence, the mechanism can be implemented in polynomial time. ■ (Proof of Theorem 5.3)

5.5 Experiments

We presented two potentially useful mechanisms, DYNAMIC DRF and CAUTIOUS LP, each with its own theoretical guarantees. Our next goal is to analyze the performance of both mechanisms on *real data*, for two natural objectives: the sum of dominant shares (the maxsum objective) and the minimum dominant share (the maxmin objective) of the agents present in the system.[1]

We compare the objective function values achieved by the two mechanisms with certain lower and upper bounds. Since both mechanisms satisfy SI, their maxsum and maxmin objective values are provably lower bounded by k/n and $1/n$, respectively, at step k.

For upper bounds, we consider the omniscient (hence unrealistic) mechanisms that maximize the objectives in an offline setting where the mechanisms have complete knowledge of future demands. These mechanisms need to guarantee an EF extension only on the real future demands rather than on all possible future demands. The comparison of CAUTIOUS LP with these offline mechanisms demonstrates the loss CAUTIOUS LP (an online mechanism) suffers due to the absence of information regarding the future demands, that is, due to its cautiousness. Because DYNAMIC DRF is not required to have an EF extension, the offline mechanisms are not theoretical upper bounds for DYNAMIC DRF, but our experiments show that they provide upper bounds in practice.

As our data we use traces of real workloads[2] on a Google compute cell, from a 7 hour period in 2011 [111]. The workload consists of tasks, where each task ran on a single machine, and consumed memory and one or more cores; the demands fit our model with two resources assuming both memory and computing power are perfectly divisible resources. For various values of n, we sampled n random positive demand vectors from the traces and analyzed the value of the two objective functions under DYNAMIC DRF and CAUTIOUS LP along with the corresponding lower and upper bounds. We averaged over 1000 such simulations to obtain data points.

Figures 5.4(a) and 5.4(b) show the maxsum values achieved by the different mechanisms, for 20 agents and 100 agents respectively. The performance of our two mechanisms is nearly identical.

Figures 5.4(c) and 5.4(d) show the maxmin values achieved for 20 agents and 100 agents, respectively. Observe that DYNAMIC DRF performs better than CAUTIOUS LP for lower values of k, but performs worse for higher values of k. Intuitively, DYNAMIC DRF allocates more resources in early stages to satisfy DPO while CAUTIOUS LP cautiously waits. This results in the superior performance of DYNAMIC DRF in initial steps but it has fewer resources available and thus lesser flexibility for optimization in later steps, resulting in inferior performance near the end. In contrast, CAUTIOUS LP is later able to make up for its loss in early steps. Encouragingly, by the last step CAUTIOUS LP achieves near optimal maxmin value. For the same reason, unlike DYNAMIC DRF the maxmin objective value for CAUTIOUS LP monotonically increases as k increases in our experiments (although it is easy to show that this is not always the case).

[1]Under a cardinal notion of utility where the dominant share of an agent is its utility, the sum of dominant shares is the utilitarian social welfare and the minimum dominant share is the egalitarian social welfare.

[2]These are the reported workloads, and may not be truthful revelations in the absence of a strategyproof mechanism.

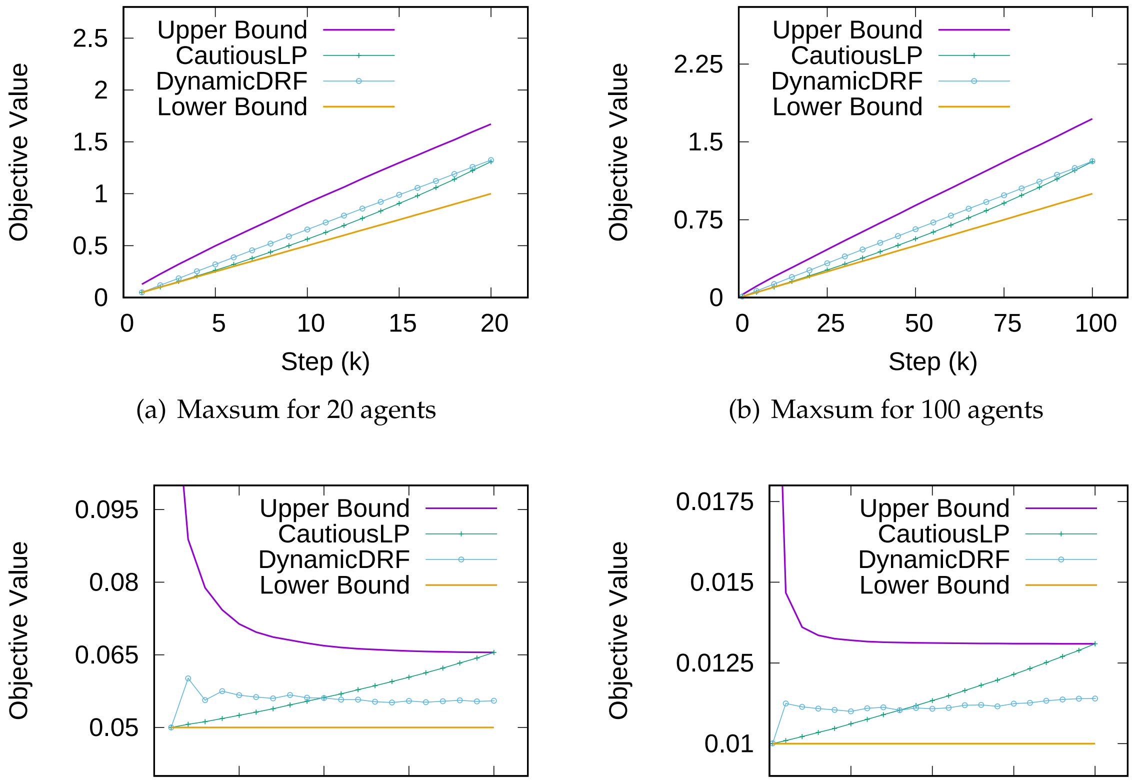

(a) Maxsum for 20 agents

(b) Maxsum for 100 agents

(c) Maxmin for 20 agents

(d) Maxmin for 100 agents

Figure 5.4: The maxsum and maxmin objectives as a function of the time step k, for $n = 20$ and $n = 100$.

5.6 Related Work

Walsh [198] proposed the problem of fair cake cutting where agents arrive, take a piece of cake, and immediately depart. The cake cutting setting deals with the allocation of a *single, heterogeneous* divisible resource; contrast with our setting, which deals with *multiple, homogeneous* divisible resources. Walsh suggested several desirable properties for cake cutting mechanisms in this setting, and showed that adaptations of classic mechanisms achieve these properties (Walsh also pointed out that allocating the whole cake to the first agent achieves the same properties). As we note in Section 5.3, his notion of *forward envy freeness* is related to our notion of *dynamic envy freeness*.

The networking community has studied the problem of fairly allocating a single homogeneous resource in a queuing model where each agent's task requires a given number of time units to be processed. In other words, in these models tasks are processed over time, but demands stay fixed, and there are no other dynamics such as agent arrivals and departures. The well-known *fair queuing* solution [73] allocates one unit per agent in successive round-robin fashion. This solution has also been analyzed by economists [152].

Previous papers on the allocation of multiple resources study a static setting. For example, Ghodsi et al. [99] proposed the *dominant resource fairness* (DRF) mechanism, which guarantees a number of desirable theoretical properties. Li and Xue [132] presented characterizations of mechanisms satisfying various desiderata while Joe-Wong et al. [116] analyzed the classic tradeoff between fairness and efficiency, both in generic frameworks that capture DRF as a special case. Parkes et al. [165] extended DRF in several ways, and in particular studied the case of indivisible tasks. Finally, DRF has also been extended to the queuing domain [100] and to incorporate job placement considerations [101], but these generalizations also use a static setting. Recently, Zahedi and Lee [211] applied the concept of Competitive Equilibrium from Equal Outcomes (CEEI) in the case of Cobb-Douglas utilities to achieve properties similar to DRF. They empirically show that these utilities are well suited for modeling user preferences over hardware resources such as cache capacity and memory bandwidth. Dolev et al. [78] defined a notion of fairness that is different from the one considered in DRF. They also proved that a fair allocation according to this new notion is always guaranteed to exist in a static setting. Gutman and Nisan [108] gave a polynomial time algorithm to find such an allocation, and also considered generalizations of DRF in a more general model of utilities.

5.7 Discussion

We have presented a new model for resource allocation with multiple resources in dynamic environments that, we believe, can spark the study of dynamic fair division more generally. The model is directly applicable to data centers, clusters, and cloud computing, where the allocation of multiple resources is a key issue, and it significantly extends the previously studied static models. That said, the model also gives rise to technical challenges that need to be tackled to capture more realistic settings.

First, our model assumes positive demands, that is, each agent requires every resource. To see how the positive demands assumption plays a role, recall that achieving EF and DPO is impossible. We established that dropping DPO leads to the trivial mechanism EQUAL SPLIT, which satisfies the remaining three properties; this is also true for possibly zero demands. When we dropped EF, we observed that the trivial mechanism DYNAMIC DICTATORSHIP satisfies SI, DPO and SP, and we subsequently suggested the improved mechanism DYNAMIC DRF that satisfies DEF in addition to SI, DPO and SP. Surprisingly though, it can be shown that neither DYNAMIC DICTATORSHIP (see Example 5.2) nor DYNAMIC DRF are SP under possibly zero demands.[3] In fact, despite significant effort, we were unable to settle the question of the existence of a mechanism that satisfies SI, DPO and SP under possibly zero demands.

[3]Under possibly zero demands, we modify DYNAMIC DICTATORSHIP and DYNAMIC DRF to continue allocating even when some resources become saturated so that they satisfy DPO.

Second, our analysis is restricted to the setting of divisible tasks, where agents value fractional quantities of their tasks. Parkes et al. [165] consider the indivisible tasks setting, where only integral quantities of an agent's task are executed, albeit in a static environment. It can be shown that even forward EF — the weakest of all EF relaxations considered in this chapter — is impossible to achieve along with DPO under indivisible tasks. It remains open to determine which relaxations of EF are feasible in dynamic resource allocation settings with indivisible tasks. While we restrict our attention to Leontief utilities, it should be noted that the desiderata we propose are well-defined in our dynamic setting with any utility function.

Third, while our model of fair division extends the classical model by introducing dynamics, and our results can directly inform the design of practical mechanisms, we do make the assumption that agents arrive over time but do not depart. In reality, agents may arrive and depart multiple times, and their preferences may also change over time (note that changing preferences can be modeled as a departure and simultaneous re-arrival with a different demand vector). Departures without re-arrivals are easy to handle; one can allocate the resources that become free in a similar way to allocations of entitlements, e.g., using DYNAMIC DRF (this scheme would clearly satisfy SI, DEF, and DPO, and it would be interesting to check whether it is also strategyproof). However, departures with re-arrivals immediately lead to daunting impossibilities. Note though that mechanisms that were designed for static settings performed well in realistic (fully dynamic) environments [99], and it is quite likely that our mechanisms — which do provide theoretical guarantees for restricted dynamic settings — would yield even better performance in reality.

Part II

Social Choice Theory

Chapter 6

Subset Selection Using Implicit Utilitarian Voting

6.1 Introduction

In this chapter, we study the classic social choice problem of aggregating the preferences of a set of voters — represented as *rankings* over a set of alternatives — into a collective decision. While traditional social choice theory takes a *normative* approach, by specifying desirable axioms that the aggregation method (also known as a *voting rule*) should satisfy [5], we take a *quantitative* approach in which one identifies and optimizes a compelling objective function.

Procaccia and Rosenschein [171] propose optimizing the *utilitarian social welfare*. Specifically, they assume that each voter has a *cardinal* utility for each possible outcome, but expresses only *ordinal* preferences (a ranking) that reflects the order of the outcomes by their cardinal utilities. The goal is to select the socially optimal outcome that maximizes the sum of the latent cardinal utilities. The performance of a voting rule — which can only access the submitted rankings, not the implicit utility functions — can then be quantified via a measure called *distortion*: the worst-case (over utility functions consistent with the reported profile of rankings) ratio between the social welfare of the optimal (welfare-maximizing) alternative, and the social welfare of the alternative selected by the voting rule. While Procaccia and Rosenschein [171] focus on analyzing the distortion of existing voting rules, Boutilier et al. [35] design voting rules that minimize distortion. In particular, they bound the worst-case distortion, and show that the distortion-minimizing (randomized) voting rule can be implemented inf polynomial time. We refer to this approach as *implicit utilitarian voting*.[1]

The work of Boutilier et al. [35] provides a good understanding of optimized aggregation of rankings from the utilitarian viewpoint — but only when a single alternative is selected by the voting rule. Indeed, this understanding does not extend to common applications that require selection of a subset of alternatives, such as choosing a com-

[1]Cf. *utilitarian voting*, which has sporadically been used to refer to both approval voting and range voting.

mittee, or selecting restaurants for the next four group lunches. In this chapter, our goal is therefore to

> *...build on the utilitarian approach to design optimal voting rules for selecting a subset of alternatives, and understand the guarantees they provide, as well as their performance in practice.*

6.1.1 Direct Real-World Implications

Research in computational social choice has frequently been justified by potential applications in multiagent systems. But recently researchers have begun to realize that, arguably, the most exciting products of this research are computer programs that help *humans* make decisions via AI-driven algorithms. One example is our aforementioned fair division website Spliddit.org [104]. In the voting space, existing examples include *Whale* (whale3.noiraudes.net/whale3/) and *Pnyx* (pnyx.dss.in.tum.de) — but these websites generally adopt the axiomatic viewpoint.

Since May 2015, some of us have been working on the design and implementation of a new not-for-profit social choice website, RoboVote.org, which is scheduled to launch in the fall of 2016. The novelty of RoboVote is that it relies on optimization-based approaches. For the case of objective votes — when a ground truth ranking of the alternatives exists (e.g., the order of different stocks by the relative change in their prices tomorrow) — RoboVote implements voting rules that pinpoint the most likely best alternative [209], or the set most likely to contain it [179]. For the case of subjective votes — the classic setting which is the focus of this chapter, with applications to everyday scenarios such as a group of friends selecting a movie to watch or a restaurant to go to — we use the results of Boutilier et al. [35] to select a single alternative. But, previously, the extension to subset selection was unavailable — this is precisely the motivation for the work described herein. Based on the results of Sections 6.4 and 6.5, we have implemented the deterministic regret minimization rule on RoboVote.

6.2 The Model

Let $[t] = \{1, \ldots, t\}$. Let A be the set of *alternatives*, and denote $m = |A|$. Let $N = [n]$ be the set of *voters*. Let $\mathcal{L} = \mathcal{L}(A)$ denote the set of rankings over the alternatives. Each voter $i \in [n]$ submits a ranking $\sigma_i \in \mathcal{L}$ over the alternatives, and which can alternatively be seen as a permutation of A. Therefore, $\sigma_i(a)$ is the position in which voter i ranks alternative a (1 is best, m is worst). Moreover, $a \succ_{\sigma_i} b$ denotes that voter i prefers alternative a to alternative b. The collection of voters' (submitted) rankings is called the *preference profile*, and denoted by $\vec{\sigma} \in \mathcal{L}^n$.

We assume the rankings are induced by comparisons between the voters' underlying utilities. For $i \in N$ and $a \in A$, let $u_i(a) \in [0, 1]$ be the utility of voter i for alternative a. As in previous papers [35, 47], we assume that the utilities are normalized such that $\sum_{a \in A} u_i(a) = 1$ for all $i \in N$. The collection of voter utilities, denoted $\vec{u}$, is called the *utility profile*. We say that utility profile $\vec{u}$ is consistent with preference profile $\vec{\sigma}$ — denoted $\vec{u} \triangleright \vec{\sigma}$ — if for all $a, b \in A$ and $i \in N$, $a \succ_{\sigma_i} b$ implies $u_i(a) \geqslant u_i(b)$.

Next we need to define the utility of a voter for a set of alternatives. For $S \subseteq A$, we define $u_i(S) = \max_{a \in S} u_i(a)$, that is, each voter derives utility for his favorite alternative in the set; this is in the same spirit as previous papers on set selection [54, 138, 144, 178]. Then, the (utilitarian) *social welfare* of S given the utility profile $\vec{u}$ is $\texttt{sw}(S, \vec{u}) = \sum_{i=1}^{n} u_i(S)$.

We are interested in voting rules that, given a preference profile, select a subset of given cardinality k.[2] Therefore, it will be useful to denote $\mathcal{A}_k = \{S \subseteq A : |S| = k\}$. In order to unify notation, we directly define a *randomized* voting rule as a function $f : \mathcal{L}^n \to \Delta(\mathcal{A}_k)$, that is, the rule is allowed to select alternatives randomly, and formally $f(\vec{\sigma})$ is a *probability distribution* over $\mathcal{A}_k$. A *deterministic* voting rule simply gives probability 1 to a specific subset.

A voting rule can only access the preference profile $\vec{\sigma}$, yet the goal is to maximize social welfare with respect to the latent utility function $\vec{u} \triangleright \vec{\sigma}$. We study two notions that quantify how well a rule achieves this goal: distortion and regret.

The *distortion* [171] of a (randomized) voting rule f on a preference profile $\vec{\sigma}$ is

$$\texttt{dist}(f, \vec{\sigma}) = \sup_{\vec{u} \triangleright \vec{\sigma}} \frac{\max_{S \in \mathcal{A}_k} \texttt{sw}(S, \vec{u})}{\mathbb{E}[\texttt{sw}(f(\vec{\sigma}), \vec{u})]}.$$

In words, it is the worst-case — over utility profiles consistent with the given preference profile — ratio between the social welfare of the best subset, and the expected social welfare of the subset selected by the voting rule. We define the distortion of a voting rule f by taking the worst case over preference profiles: $\texttt{dist}(f) = \max_{\vec{\sigma} \in \mathcal{L}^n} \texttt{dist}(f, \vec{\sigma})$.

The second measure is *regret*. While it has not been studied as part of the agenda of implicit utilitarian voting, it has been explored in other social choice settings, especially partial preferences [136]; similar measures have been extensively studied in decision theory and machine learning [27, 43]. The *regret* of a (randomized) voting rule f on a preference profile $\vec{\sigma}$ is given by

$$\texttt{reg}(f, \vec{\sigma}) = \frac{1}{n} \cdot \sup_{\vec{u} \triangleright \vec{\sigma}} \left(\max_{S \in \mathcal{A}_k} \texttt{sw}(S, \vec{u}) - \mathbb{E}[\texttt{sw}(f(\vec{\sigma}), \vec{u})] \right).$$

As before, define the regret of a rule f to be $\texttt{reg}(f) = \max_{\vec{\sigma} \in \mathcal{L}^n} \texttt{reg}(f, \vec{\sigma})$. We divide by n because the total (worst-case) regret of any voting rule f is provably linear in n (so this is *per vote* regret). Note that distortion is a multiplicative measure of loss, whereas regret is its additive version.

6.3 Worst-Case Bounds

In this section we provide bounds on worst-case distortion and regret, for both deterministic and randomized voting rules. Boutilier et al. [35] show that for selecting a single winner ($k = 1$), we can achieve $O(\sqrt{m} \cdot \log^* m)$ distortion using a randomized rule,

[2]Formally, this is a special case of *social choice correspondences* with fixed output cardinality [46].

where $\log^* m$ is the iterated logarithm of m (the number of alternatives). This bound is asymptotically almost tight: they also show that the worst-case distortion is always $\Omega(\sqrt{m})$.

For a large k, though, one can hope for a better bound. Clearly, when $k = m$ there is only one voting rule (which selects every alternative), and its distortion is 1. More generally, it is easy to show that the voting rule f that selects a subset from $\mathcal{A}_k$ uniformly at random has $\texttt{dist}(f) \leqslant m/k$. However, since we can already achieve $O(\sqrt{m} \cdot \log^* m)$ distortion for $k = 1$, a bound of m/k provides an improvement only for $k = \Omega(\sqrt{m}/\log^* m)$. Can we achieve better distortion for smaller values of k as well? It is not even clear whether the optimal worst-case distortion should monotonically decrease in k, because as our flexibility grows with k, so does the flexibility of the welfare-maximizing solution. In fact, a part of our main result shows that the worst-case distortion remains $\Omega(\sqrt{m})$ for all values of k up to $\Theta(\sqrt{m})$.

Theorem 6.1. *Let $m = |A|$, and let k be the number of alternatives to be selected.*

1. ***Distortion, deterministic rules:*** *There exists a deterministic voting rule f^* with $\texttt{dist}(f^*) \leqslant 1 + m\,(m-k)/k$. Moreover, for every deterministic voting rule f,*

$$\texttt{dist}(f) \geqslant \begin{cases} 1 + \frac{m(m-3k)}{6k} & \text{if } k \leqslant \frac{m}{9}, \\ 1 + m & \text{if } \frac{m}{9} < k \leqslant \frac{m}{2}, \\ 1 + \frac{m(m-k)}{k} & \text{otherwise.} \end{cases}$$

   These bounds are tight up to a constant factor of 8.

2. ***Distortion, randomized rules:*** *There exists a randomized voting rule f^* such that*

$$\texttt{dist}(f^*) \leqslant \begin{cases} 2\sqrt{m \cdot H_m} & \text{if } k \leqslant \frac{2 \cdot m \cdot H_m}{m + H_m}, \\ 4\sqrt{m \cdot k} & \text{if } \frac{2 \cdot m \cdot H_m}{m + H_m} < k \leqslant \left(\frac{m}{4}\right)^{\frac{1}{3}}, \\ \frac{m}{k} & \text{otherwise,} \end{cases}$$

   where $H_m = \Theta(\log m)$ is the m^{th} harmonic number. Moreover, for every randomized voting rule f,

$$\texttt{dist}(f) \geqslant \begin{cases} \frac{\sqrt{m}}{2} & \text{if } k \leqslant \frac{m \cdot (\sqrt{m} - 1)}{m - 1} \approx \sqrt{m}, \\ \frac{m}{k + m/k} & \text{otherwise.} \end{cases}$$

   These bounds are tight up to a factor of $6.35 \cdot m^{1/6}$.

3. ***Regret, deterministic rules:*** *There exists a deterministic voting rule f^* such that*

$$\texttt{reg}(f^*) \leqslant \begin{cases} \frac{1}{2} & \text{if } k \leqslant \frac{m}{2}, \\ 1 - \frac{k}{m} & \text{otherwise,} \end{cases}$$

   and this upper bound is completely tight.

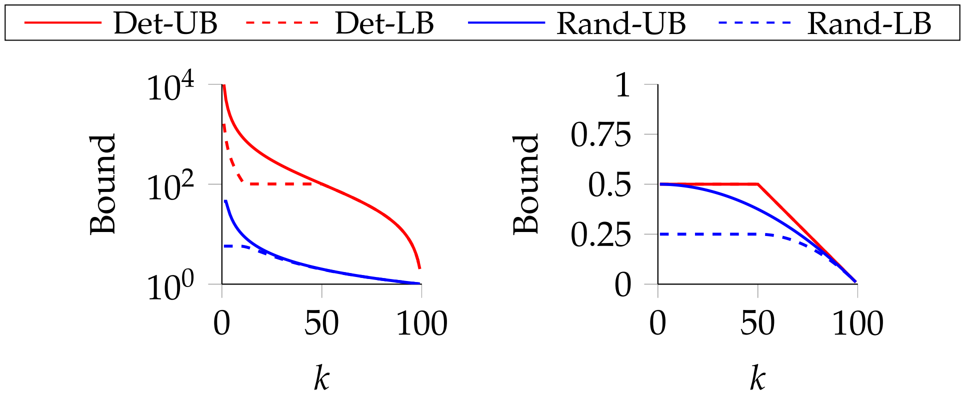

(a) Distortion (b) Regret

Figure 6.1: The upper and lower bounds on worst-case distortion and regret for $m = 100$.

4. ***Regret, randomized rules:*** *There exists a randomized voting rule f^* such that*

$$\mathtt{reg}(f^*) \leqslant \frac{1}{2} \cdot \left(1 - \frac{k^2}{m^2}\right).$$

Moreover, for every randomized voting rule f,

$$\mathtt{reg}(f) \geqslant \begin{cases} \frac{1}{4} & \text{if } k \leqslant m/2 \\ \frac{1}{2} \cdot \frac{k}{m}\left(1 - \frac{k}{m}\right) & \text{otherwise.} \end{cases}$$

These bounds are tight up to a constant factor of 2.

All the upper bounds above can be achieved via polynomial-time algorithms.

The bounds presented above are simplified forms of the exact bounds that we derive. Figure 6.1 shows our exact bounds for $m = 100$.[3]

Before we dive into the proof, we simplify the formulae for the distortion and regret of deterministic voting rules.

Definition 6.1 (Comparing Sets)**.** Given a ranking $\sigma \in \mathcal{L}$ and an alternative $a \in A$, recall that $\sigma(a)$ denotes the position of a in σ. More generally, for a set $S \subseteq A$ let $\sigma(S) = \min_{a \in S} \sigma(a)$. For sets $S, T \subseteq A$, we say $T \succ_\sigma S$ if $\sigma(T) < \sigma(S)$, i.e., if there exists an alternative in T that is preferred to every alternative in S in σ.

Definition 6.2 (Plurality Score)**.** The plurality score of an alternative $a \in A$ in a preference profile $\vec{\sigma}$ is the number of votes in which a appears first, i.e.,

$$\mathtt{plu}(a, \vec{\sigma}) = \sum_{i=1}^{n} \mathbb{I}[\sigma_i(1) = a],$$

[3]The second upper bound in part 2 of Theorem 6.1 (which increases with k) does not play a role unless m is very large.

where $\mathbb{I}$ is the indicator function. More generally, we define the plurality score of a set $S \subseteq A$ to be the number of votes in which an alternative in S is ranked first, i.e.,

$$\texttt{plu}(S, \vec{\sigma}) = \sum_{i=1}^{n} \mathbb{I}[\sigma_i(S) = 1] = \sum_{a \in S} \texttt{plu}(a, \vec{\sigma}).$$

Lemma 6.1. *For a deterministic voting rule f and a preference profile $\vec{\sigma}$, the regret of f on $\vec{\sigma}$ is given by*

$$\texttt{reg}(f, \vec{\sigma}) = \max_{S \in \mathcal{A}_k} \frac{1}{n} \cdot \sum_{i=1}^{n} \frac{\mathbb{I}[S \succ_{\sigma_i} f(\vec{\sigma})]}{\sigma_i(S)}, \tag{6.1}$$

and the distortion of f on $\vec{\sigma}$ is given by

$$\texttt{dist}(f, \vec{\sigma}) = 1 + m \cdot \frac{n \cdot \texttt{reg}(f, \vec{\sigma})}{\texttt{plu}(f(\vec{\sigma}), \vec{\sigma})}. \tag{6.2}$$

Proof. First, note that $\texttt{reg}(f, \vec{\sigma})$ and $\texttt{dist}(f, \vec{\sigma})$ can be rewritten as follows.

$$\texttt{reg}(f, \vec{\sigma}) = \frac{1}{n} \cdot \sup_{\vec{u} \rhd \vec{\sigma}} \left[\max_{S \in \mathcal{A}_k} \texttt{sw}(S, \vec{u}) - \texttt{sw}(f(\vec{\sigma}), \vec{u}) \right] = \frac{1}{n} \cdot \left[\max_{S \in \mathcal{A}_k} \sup_{\vec{u} \rhd \vec{\sigma}} \texttt{sw}(S, \vec{u}) - \texttt{sw}(f(\vec{\sigma}), \vec{u}) \right], \tag{6.3}$$

and similarly,

$$\texttt{dist}(f, \vec{\sigma}) = \sup_{\vec{u} \rhd \vec{\sigma}} \frac{\max_{S \in \mathcal{A}_k} \texttt{sw}(S, \vec{u})}{\texttt{sw}(f(\vec{\sigma}), \vec{u})} = \max_{S \in \mathcal{A}_k} \sup_{\vec{u} \rhd \vec{\sigma}} \frac{\texttt{sw}(S, \vec{u})}{\texttt{sw}(f(\vec{\sigma}), \vec{u})}. \tag{6.4}$$

If $S = f(\vec{\sigma})$, then the regret term is 0 and the distortion term is 1. Fix $S \in \mathcal{A}_k \setminus \{f(\vec{\sigma})\}$. To maximize the regret (resp. distortion) term, we want to find the utility profile $\vec{u}$ that maximizes the difference (resp. ratio) of $\texttt{sw}(S, \vec{u})$ and $\texttt{sw}(f(\vec{\sigma}), \vec{u})$ subject to $\vec{u} \rhd \vec{\sigma}$. Let us construct a specific utility profile $\vec{u}^*$ where the utility function of voter i is given as follows. For each voter i,

1. If $\sigma_i(f(\vec{\sigma})) = 1$, let $u_i(a) = 1/m$ for all $a \in A$.
2. If $S \succ_i f(\vec{\sigma})$, let $u_i(a) = 1/\sigma_i(S)$ if $\sigma_i(a) \in [\sigma_i(S)]$, and $u_i(a) = 0$ otherwise.
3. If $S \not\succ_i f(\vec{\sigma})$ and $\sigma_i(f(\vec{\sigma})) \neq 1$, let $u_i(a) = 1$ if $\sigma_i(a) = 1$, and $u_i(a) = 0$ otherwise.

First, it is easy to check that $\vec{u}^* \rhd \vec{\sigma}$. Also, note that this utility profile maximizes $u_i(S) - u_i(f(\vec{\sigma}))$ subject to $u_i \rhd \sigma_i$, for each voter i in each of the three cases above. Hence, it maximizes the regret term $\texttt{sw}(S, \vec{u}) - \texttt{sw}(f(\vec{\sigma}), \vec{u})$. Further, we have

$$\begin{aligned} \texttt{sw}(f(\vec{\sigma}), \vec{u}^*) &= \frac{1}{m} \cdot \texttt{plu}(f(\vec{\sigma}), \vec{\sigma}), \\ \texttt{sw}(S, \vec{u}^*) &= \frac{1}{m} \cdot \texttt{plu}(f(\vec{\sigma}), \vec{\sigma}) + \sum_{i=1}^{n} \frac{\mathbb{I}[S \succ_i f(\vec{\sigma})]}{\sigma_i(S)}. \end{aligned}$$

This immediately gives us Equation (6.1) for the regret of f on $\vec{\sigma}$. Now, the distortion of f on $\vec{\sigma}$ under the utility profile $\vec{u}^*$ is

$$\frac{\mathtt{sw}(S,\vec{u}^*)}{\mathtt{sw}(f(\vec{\sigma}),\vec{u}^*)} = 1 + m \cdot \frac{\sum_{i=1}^{n} \mathbb{I}[S \succ_i f(\vec{\sigma})]/\sigma_i(S)}{\mathtt{plu}(f(\vec{\sigma}),\vec{\sigma})} = 1 + m \cdot \frac{n \cdot \mathtt{reg}(f,\vec{\sigma})}{\mathtt{plu}(f(\vec{\sigma}),\vec{\sigma})} \geqslant 1. \tag{6.5}$$

Finally, take a utility profile $\vec{u}$ satisfying $\vec{u} \triangleright \vec{\sigma}$. We want to show that the distortion under $\vec{u}$, i.e., $\mathtt{sw}(S,\vec{u})/\mathtt{sw}(f(\vec{\sigma}),\vec{u})$ is no more than the distortion under $\vec{u}^*$. Note that $f(\vec{\sigma})$ has the least possible welfare under $\vec{u}^*$. Hence, $\mathtt{sw}(f(\vec{\sigma}),\vec{u}) \geqslant \mathtt{sw}(f(\vec{\sigma}),\vec{u}^*)$. To achieve a greater distortion, we must have $\mathtt{sw}(S,\vec{u}) > \mathtt{sw}(S,\vec{u}^*)$, i.e., the voters must assign a greater utility to S under $\vec{u}$ than under $\vec{u}^*$.

Let us revisit the three cases in the construction of $\vec{u}^*$. All the voters covered under case 2 already assign S the highest possible utility. For all the other voters, the top-ranked alternative of $f(\vec{\sigma})$ is at least as high as the top-ranked alternative of S. Hence, an increase of δ in the utility for S would require an increase of at least δ in the utility for $f(\vec{\sigma})$.

In other words, for any $\delta > 0$, $\mathtt{sw}(S,\vec{u}) = \mathtt{sw}(S,\vec{u}^*) + \delta$ implies $\mathtt{sw}(f(\vec{\sigma}),\vec{u}) \geqslant \mathtt{sw}(f(\vec{\sigma}),\vec{u}^*) + \delta$. Finally,

$$\frac{\mathtt{sw}(S,\vec{u}^*)}{\mathtt{sw}(f(\vec{\sigma}),\vec{u}^*)} \geqslant 1 \Rightarrow \frac{\mathtt{sw}(S,\vec{u})}{\mathtt{sw}(f(\vec{\sigma}),\vec{u})} \leqslant \frac{\mathtt{sw}(S,\vec{u}^*)+\delta}{\mathtt{sw}(f(\vec{\sigma}),\vec{u}^*)+\delta} \leqslant \frac{\mathtt{sw}(S,\vec{u}^*)}{\mathtt{sw}(f(\vec{\sigma}),\vec{u}^*)}.$$

Hence, the worst-case distortion is indeed $1 + m \cdot \mathtt{reg}(f,\vec{\sigma})/\mathtt{plu}(f(\vec{\sigma}),\vec{\sigma})$, as required. Note that in finding the worst-case distortion, the distortion of 1 achieved with $S = f(\vec{\sigma})$ is ignored because the distortion achieved by every $S \neq f(\vec{\sigma})$ is at least 1. ■

One interesting consequence of Lemma 6.1 is that selecting a set of alternatives, none of which appear at the top position in any vote, results in an unbounded distortion. Hence, the rule that optimizes distortion would never select such a set. We are now ready for the proof of our main result.

Proof of Theorem 6.1. Below, we provide a proof for the upper and the lower bound in each of the four cases.

Distortion, deterministic rules:

Upper bound: Inspired by the denominator $\mathtt{plu}(f(\vec{\sigma}),\vec{\sigma}) = \sum_{a \in f(\vec{\sigma})} \mathtt{plu}(a,\vec{\sigma})$ in the formula for the distortion of f on $\vec{\sigma}$ from Lemma 6.1, let us analyze the rule f^* that selects the k alternatives with the highest plurality scores. We show that f^* achieves the required upper bound on the optimal distortion.

Since the sum of plurality scores of all the alternatives is n, the sum of top k plurality scores is at least $n \cdot k/m$. Hence, $\mathtt{plu}(f^*(\vec{\sigma}),\vec{\sigma}) \geqslant n \cdot k/m$. Next, for $S \in \mathcal{A}_k \setminus \{f^*(\vec{\sigma})\}$, note that the number of voters i for which $S \succ_i f^*(\vec{\sigma})$ is at most $n - \mathtt{plu}(f^*(\vec{\sigma}),\vec{\sigma})$. Using Lemma 6.1, it follows that the distortion of f^* on $\vec{\sigma}$ is at most

$$1 + m \cdot \frac{n - \mathtt{plu}(f^*(\vec{\sigma}),\vec{\sigma})}{\mathtt{plu}(f^*(\vec{\sigma}),\vec{\sigma})} = 1 + m \cdot \left(\frac{n}{\mathtt{plu}(f^*(\vec{\sigma}),\vec{\sigma})} - 1\right) \leqslant 1 + m \cdot \left(\frac{m}{k} - 1\right),$$

which is the required upper bound.

Lower bound: Next, we establish three different lower bounds on the distortion of deterministic rules.

1. For $k \leqslant m/6$, $\texttt{dist}(f) \geqslant 1 + m \cdot (m-3k)/(6k)$ for every deterministic rule f.
2. For $k \leqslant m/2$, $\texttt{dist}(f) \geqslant 1 + m$ for every deterministic rule f.
3. For $k \geqslant m/2$, $\texttt{dist}(f) \geqslant 1 + m \cdot (m-k)/k$ for every deterministic rule f.

It is easy to check that for $k \leqslant m/9$, the first bound is the strongest; for $k \in [m/9, m/2]$, the second bound is the strongest; and, for $k \geqslant m/2$, the third bound is the strongest. Hence, the optimal combination of these three bounds gives us the desired result.

Let the set of alternatives be $A = \{a_1, \ldots, a_m\}$. Let us begin by proving the first bound. Fix a value of $k \leqslant m/6$, and partition A into two sets: $X = \{a_1, \ldots, a_{m-3k}\}$ and $Y = A \setminus X$. Let $Y = \{b_1, \ldots, b_{3k}\}$. Note that $m \geqslant 6k$ implies $|X| \geqslant |Y|$. Construct a ranked profile $\vec{\sigma}$ as follows.

- For each alternative $a_j \in X$ (thus $j \leqslant m-3k$), let a_j be ranked first in the votes of voters $i \in [(j-1) \cdot n/(m-3k)+1, j \cdot n/(m-3k)]$. That is, we partition the set of voters into $|X| = m-3k$ contiguous blocks, and have each alternative of X ranked first in one of the blocks.
- For each alternative $b_j \in Y$ (thus $j \leqslant 3k$), let b_j be ranked second in the votes of voters $i \in [(j-1) \cdot n/3k+1, j \cdot n/3k]$. That is, we partition the set of voters into $|Y| = 3k$ contiguous blocks, and have each alternative of Y ranked second in one of the blocks.

Since we chose the blocks of voters to be contiguous in both cases, it follows that for every $a_j \in X$, the set of voters ranking a_j first can have at most two distinct alternatives in Y in their second position. Take a deterministic rule f for selecting a set of k alternatives. Let $|f(\vec{\sigma}) \cap X| = t \leqslant k$. Then, we have $\texttt{plu}(f(\vec{\sigma}), \vec{\sigma}) = t \cdot n/(m-3k)$.

Consider the voters who rank an alternative of $f(\vec{\sigma})$ first. Let Y' denote the set of alternatives appearing in the second position in the votes of such voters. From the argument above, we have $|Y'| \leqslant 2t \leqslant 2k$. Hence, $|Y \setminus Y'| \geqslant 3k - 2k = k$. Choose an arbitrary set $S \subseteq Y \setminus Y'$ such that $|S| = k$. Now, there are $(n/3k) \cdot k = n/3$ voters that rank an alternative in S in their second position. Hence,

$$\begin{aligned}\texttt{dist}(f) \geqslant \texttt{dist}(f, \vec{\sigma}) \geqslant 1 + m \cdot \frac{\sum_{i=1}^{n} \mathbb{I}[S \succ_i f(\vec{\sigma})]/\sigma_i(S)}{\texttt{plu}(f(\vec{\sigma}), \vec{\sigma})} &\geqslant 1 + m \cdot \frac{(n/3) \cdot (1/2)}{t \cdot n/(m-3k)} \\ &\geqslant 1 + m \cdot \frac{m-3k}{6k},\end{aligned}$$

where the second transition uses Lemma 6.1, and the final transition uses $t \leqslant k$.

For the second and the third lower bound, we simply construct a profile $\vec{\sigma}$ in which each alternative in A appears first in n/m votes, and the remaining positions in the votes are filled arbitrarily. Fix a deterministic rule f with $|f(\vec{\sigma})| = k$. Note that $\texttt{plu}(f(\vec{\sigma}), \vec{\sigma}) = (n/m) \cdot k$.

If $k \leqslant m/2$, choose a set $S \subseteq A \setminus f(\vec{\sigma})$ such that $|S| = k$. Note that $S \cap f(\vec{\sigma}) = \emptyset$, and an alternative in S is ranked first in $(n/m) \cdot k$ votes. Hence, by Lemma 6.1,

$$\texttt{dist}(f) \geqslant \texttt{dist}(f, \vec{\sigma}) \geqslant 1 + m \cdot \frac{\sum_{i=1}^{n} \mathbb{I}[S \succ_i f(\vec{\sigma})]/\sigma_i(S)}{\texttt{plu}(f(\vec{\sigma}), \vec{\sigma})} \geqslant 1 + m \cdot \frac{(n/m) \cdot k}{(n/m) \cdot k} = 1 + m.$$

If $k > m/2$, choose $S \supseteq A \setminus f(\vec{\sigma})$ such that $|S| = k$. In this case, an alternative in $S \setminus f(\vec{\sigma})$ is ranked first in $(n/m) \cdot (m-k)$ votes. Hence, by Lemma 6.1,

$$\begin{aligned}\texttt{dist}(f) \geqslant \texttt{dist}(f, \vec{\sigma}) \geqslant 1 + m \cdot \frac{\sum_{i=1}^{n} \mathbb{I}[S \succ_i f(\vec{\sigma})]/\sigma_i(S)}{\texttt{plu}(f(\vec{\sigma}), \vec{\sigma})} &\geqslant 1 + m \cdot \frac{(n/m) \cdot (m-k)}{(n/m) \cdot k} \\ &= 1 + m \cdot \frac{m-k}{k},\end{aligned}$$

as required.

Gap between upper and lower bounds: Note that the gap between the upper and the lower bounds is $\max(G_1, G_2, G_3)$, where

$$G_1 = \max_{k \in [1, \frac{m}{9}]} \frac{1+m(m-k)/k}{1+m(m-3k)/(6k)}, \; G_2 = \max_{k \in [\frac{m}{9}, \frac{m}{2}]} \frac{1+m(m-k)/k}{1+m}, \text{ and } G_3 = 1.$$

For G_1, it can be checked that the ratio of the upper and lower bounds is an increasing function of k for $m \geqslant 3$. Hence, the maximum is achieved at $k = m/9$, and is equal to $1 + 8 \cdot m/(1+m) \leqslant 8$.

For G_2, the ratio of the upper and the lower bounds is clearly a decreasing function of k. Hence, the maximum is achieved at $k = m/9$, and is equal to $G_1 \leqslant 8$.

Hence, the upper and the lower bounds are tight up to a constant factor of 8.

Distortion, randomized rules:

Upper bound: In our opinion, the proof of this part is the most non-trivial. It uses a construction that builds on the one used by Boutilier et al. [35] for $k = 1$, but requires additional tools and introduces novel techniques. As mentioned at the beginning of this section, choosing a set uniformly at random from $\mathcal{A}_k$ (under which the marginal probability of every alternative being chosen is k/m) has distortion at most m/k. However, this approach does not work well if some alternatives are significantly better than others.

In that case, one may wish to choose the alternatives with probabilities proportional to their "quality". For $a \in A$, let us define its quality by its harmonic score $\texttt{har}(a, \vec{\sigma}) = \sum_{i \in [n]} 1/\sigma_i(a)$. Then, we wish to choose alternative a with marginal "probability" $k \cdot \texttt{har}(a, \vec{\sigma}) / \sum_{b \in A} \texttt{har}(b, \vec{\sigma})$. Note that this quantity may be greater than 1. Moreover, this approach fails when all sets are almost equally good. Hence, we employ a combination of the two approaches.

Fix $0 \leqslant \alpha \leqslant 1$, and for an alternative $a \in A$ define

$$p_a = \alpha \cdot \frac{k}{m} + (1-\alpha) \cdot \frac{k \cdot \texttt{har}(a, \vec{\sigma})}{\sum_{b \in A} \texttt{har}(b, \vec{\sigma})}. \tag{6.6}$$

Using the bihierarchy extension [45] of the Birkhoff-von Neumann theorem [25, 197], we can show that there exists a distribution over $\mathcal{A}_k$ under which the marginal probabilities of selected alternatives are consistent with Equation (6.6) if and only if

$$\forall a \in A, 0 \leqslant p_a \leqslant 1 \quad \text{and} \quad \sum_{a \in A} p_a = k.$$

Note that $p_a \geqslant 0$ and $\sum_{a \in A} p_a = k$ hold by definition. The constraint $p_a \leqslant 1$ will be applied later to restrict the feasible values of α.

For now, suppose such a distribution D exists. Consider a preference profile $\vec{\sigma}$ and a utility profile $\vec{u} \triangleright \vec{\sigma}$. Let $S^* \in \arg\max_{S \in \mathcal{A}_k} \mathtt{sw}(S, \vec{u})$. Define

$$X = \sqrt{\frac{H_m}{m} \cdot \frac{\alpha}{1-\alpha}},$$

where $H_m = \sum_{t=1}^{m} 1/t$ is the m^{th} harmonic number. Note that $\sum_{a \in A} \mathtt{har}(a, \vec{\sigma}) = n \cdot H_m$. Now, consider two cases.

Case 1: Suppose $\mathtt{sw}(S^*, \vec{u}) \leqslant n \cdot X$. Then,

$$\begin{aligned}\mathbb{E}_{S \sim D}[\mathtt{sw}(S, \vec{u})] &= \sum_{S \in \mathcal{A}_k} \Pr_D[S] \cdot \left(\sum_{i=1}^{n} \max_{a \in S} u_i(a) \right) \geqslant \sum_{i=1}^{n} \left(\sum_{S \in \mathcal{A}_k} \Pr_D[S] \cdot \frac{\sum_{a \in S} u_i(a)}{k} \right) \\ &= \frac{1}{k} \sum_{i=1}^{n} \sum_{a \in A} u_i(a) \cdot \Pr_{S \sim D}[a \in S] \geqslant \frac{1}{k} \sum_{i=1}^{n} \sum_{a \in A} u_i(a) \cdot \alpha \cdot \frac{k}{m} = \alpha \cdot \frac{n}{m}.\end{aligned}$$

Hence, the distortion is

$$\frac{\mathtt{sw}(S^*, \vec{u})}{\mathbb{E}_{S \sim D}[\mathtt{sw}(S, \vec{u})]} \leqslant \frac{n \cdot X}{\alpha \cdot n/m} = \frac{X \cdot m}{\alpha} = \sqrt{\frac{m \cdot H_m}{\alpha \cdot (1-\alpha)}}.$$

Case 2: Suppose $\mathtt{sw}(S^*, \vec{u}) > n \cdot X$. Then, for each alternative $a \in S^*$, let N_a denote the subset of voters who rank a above any other alternative of S^*, i.e.,

$$N_a = \{i \in [n] : \forall b \in S^* \setminus \{a\},\ a \succ_{\sigma_i} b, \}.$$

Let $\mathtt{sw}_{N_a}(S, \vec{u})$ denote the welfare of the voters in N_a for the set of alternatives S under the utility profile $\vec{u}$. Let T_a denote the total utility that agents in N_a have for alternative a, i.e., $T_a = \sum_{i \in N_a} u_i(a)$. It can be shown (although it is nontrivial) that $\mathtt{har}(a, \vec{\sigma}) \geqslant T_a$ for all $a \in A$. Because $\{N_a\}_{a \in S^*}$ is a partition of the set of voters, we have

$$\begin{aligned}\mathbb{E}_{S \sim D}[\mathtt{sw}(S, \vec{u})] &= \mathbb{E}_{S \sim D}\left[\sum_{a \in S^*} \mathtt{sw}_{N_a}(S, \vec{u}) \right] \geqslant \sum_{a \in S^*} T_a \cdot \Pr_{S \sim D}[a \in S] \\ &\geqslant \sum_{a \in S^*} T_a \cdot (1-\alpha) \cdot \frac{k \cdot \mathtt{har}(a, \vec{\sigma})}{\sum_{b \in A} \mathtt{har}(b, \vec{\sigma})} \geqslant \frac{(1-\alpha) \cdot k}{n \cdot H_m} \cdot \sum_{a \in S^*} (T_a)^2\end{aligned}$$

$$\geqslant \frac{1-\alpha}{n \cdot H_m} \cdot \left(\sum_{a \in S^*} T_a\right)^2 = \frac{1-\alpha}{n \cdot H_m} \cdot (\texttt{sw}(S^*, \vec{u}))^2.$$

Here, the fourth transition uses $\texttt{har}(a, \vec{\sigma}) \geqslant T_a$, the fifth transition uses the power-mean inequality, and the final transition uses $\texttt{sw}(S^*, \vec{u}) = \sum_{a \in S^*} T_a$. Now, the distortion is

$$\frac{\texttt{sw}(S^*, \vec{u})}{\mathbb{E}_{S \sim D}[\texttt{sw}(S, \vec{u})]} \leqslant \frac{n \cdot H_m}{(1-\alpha) \cdot \texttt{sw}(S^*, \vec{u})} < \sqrt{\frac{m \cdot H_m}{\alpha \cdot (1-\alpha)}},$$

where the final transition uses our assumption $\texttt{sw}(S^*, \vec{u}) > n \cdot X$ along with the definition of X.

Combined analysis: In both cases, the distortion is at most $\sqrt{mH_m/(\alpha(1-\alpha))}$. The final step involves choosing the optimal value of α by minimizing this quantity subject to our constraints: $p_a \leqslant 1$ for all $a \in A$. This translates to

$$\alpha \cdot \frac{k}{m} + (1-\alpha) \cdot \frac{k \cdot \texttt{har}(a, \vec{\sigma})}{\sum_{b \in A} \texttt{har}(b, \vec{\sigma})} \leqslant 1, \forall a \in A.$$

Note that $\texttt{har}(a, \vec{\sigma}) \leqslant n$, and $\sum_{b \in A} \texttt{har}(b, \vec{\sigma}) = n \cdot H_m$. Hence, we can safely replace these constraints by the following constraint:

$$\alpha \cdot \frac{k}{m} + (1-\alpha) \cdot \frac{k}{H_m} \leqslant 1.$$

Minimizing the value of $\sqrt{mH_m/(\alpha(1-\alpha))}$ subject to this constraint, we get that:

- For $k \leqslant 2mH_m/(m+H_m)$, the optimal distortion is achieved at $\alpha = 1/2$, and is equal to $2\sqrt{m \cdot H_m}$;
- For $k > 2mH_m/(m+H_m)$, the optimal distortion is achieved at $\alpha = m \cdot (k - H_m)/(k \cdot (m - H_m))$, and is equal to

$$\frac{k \cdot (m - H_m)}{\sqrt{(m-k) \cdot (k - H_m)}}.$$

  Note that $H_m \leqslant (7/8) \cdot k$. Further, let $k \leqslant m/2$ (as this bound will anyway be replaced by a better bound for $k > m/2$). Hence, $m - k \geqslant m/2$. Substituting these, we get that the optimal distortion in this case is at most $4\sqrt{m \cdot k}$.

Finally, recall that we have a universal upper bound of m/k, achieved by choosing from $\mathcal{A}_k$ uniformly at random. One can check that $m/k < 4\sqrt{m \cdot k}$ for $k > (m/4)^{1/3}$. Hence, combining the upper bounds from the analysis above with this universal upper bound gives us the desired result.

Lower bound: Fix a set of alternatives $T = \{a_1, \ldots, a_t\}$, where $t \geqslant k$ (the exact value of t will be determined later). Partition the set of voters into t buckets; each bucket i, denoted by N_i, consists of n/t voters. Construct a ranked profile $\vec{\sigma}$ in which for all

$i \in [t]$, all voters in bucket N_i rank alternative a_i first, and the remaining alternatives arbitrarily.

Let us analyze the output of a randomized rule f on this profile. For each alternative $a \in A$, define $p_a = \Pr_{S \sim f(\vec{\sigma})}[a \in S]$. Let $X \subseteq [t]$ be the indices corresponding to the lowest k values in the sequence $(p_{a_i})_{i \in [t]}$; in other words, let X be such that $\{a_i : i \in X\}$ is the set of k alternatives from T with the lowest values of p. Now, construct a utility profile $\vec{u}$ as follows.

1. For each $i \in [t] \setminus X$, every voter in bucket N_i has utility $1/m$ for each alternative.
2. For each $i \in X$, every voter in bucket N_i has utility 1 for the alternative it ranks first (i.e., a_i), and utility 0 for the remaining alternatives.

Note that this construction is a modification of the one used in the lower bound for the deterministic case; instead of letting voters have high utility for alternatives that are *not selected*, here we let voters have high utility for alternatives that are *selected with the lowest probabilities*.

First, note that $\vec{u} \triangleright \vec{\sigma}$. Further, under $\vec{u}$ the optimal set of k alternatives is clearly $\{a_i : i \in X\}$, and its corresponding welfare is

$$\left(n - \frac{n}{t} \cdot k\right) \cdot \frac{1}{m} + \frac{n}{t} \cdot k \cdot 1,$$

because it provides utility $1/m$ to every voter in bucket N_i for $i \in [t] \setminus X$, and utility 1 to every voter in bucket N_i for $i \in X$. In contrast, the expected welfare under f is

$$\left(n - \frac{n}{t} \cdot k\right) \cdot \frac{1}{m} + \frac{n}{t} \cdot \sum_{i \in X} p_{a_i}. \tag{6.7}$$

Next, note that $\sum_{a \in T} p_a \leqslant \sum_{a \in A} p_a = k$, where the last equality follows because f always returns a set of size k. Hence, the sum of the lowest k values from $\{p_a : a \in T\}$ (i.e., $\sum_{i \in X} p_{a_i}$, by the definition of X) is at most $(k/t) \cdot k$. Substituting this in Equation (6.7), we obtain that the worst-case distortion is bounded from below by

$$\frac{\left(n - \frac{n}{t} \cdot k\right) \cdot \frac{1}{m} + \frac{n}{t} \cdot k}{\left(n - \frac{n}{t} \cdot k\right) \cdot \frac{1}{m} + \frac{n}{t} \cdot \frac{k^2}{t}} = \frac{\left(1 - \frac{k}{t}\right) \cdot \frac{1}{m} + \frac{k}{t}}{\left(1 - \frac{k}{t}\right) \cdot \frac{1}{m} + \left(\frac{k}{t}\right)^2}.$$

Finally, minimizing this with respect to t, we get that

$$\texttt{dist}(f) \geqslant \texttt{dist}(f, \vec{\sigma}) \geqslant \begin{cases} \frac{(\sqrt{m}+1)^2}{1+2\sqrt{m}} & \text{if } k \leqslant \frac{m(\sqrt{m}-1)}{m-1}, \\ \frac{m+k(m-1)}{m+k(k-1)} & \text{otherwise.} \end{cases}$$

Finally, note that

$$\frac{(\sqrt{m}+1)^2}{1+2\sqrt{m}} \geqslant \frac{(\sqrt{m}+1)^2}{2 \cdot (1+\sqrt{m})} = \frac{\sqrt{m}+1}{2} \geqslant \frac{\sqrt{m}}{2},$$

and

$$\frac{m+k(m-1)}{m+k(k-1)} = \frac{m\cdot k+m-k}{k^2+m-k} \geqslant \frac{m\cdot k}{k^2+m} = \frac{m}{k+m/k},$$

which are the required bounds.

Gap between upper and lower bounds: In this case, the gap between the upper and the lower bounds is $\max(G_1, G_2, G_3, G_4)$, where

$$\begin{aligned}
G_1 &= \max_{k\in\left[1,\frac{2mH_m}{m+H_m}\right]} \frac{2\sqrt{m\cdot H_m}}{\sqrt{m}/2} \leqslant 4\sqrt{H_m},\\
G_2 &= \max_{k\in\left[\frac{2mH_m}{m+H_m},\left(\frac{m}{4}\right)^{\frac{1}{3}}\right]} \frac{4\sqrt{m\cdot k}}{\sqrt{m}/2} \leqslant \max_{k\in\left[\frac{2mH_m}{m+H_m},\left(\frac{m}{4}\right)^{\frac{1}{3}}\right]} 8\cdot\sqrt{k} \leqslant 2^{8/3}\cdot m^{1/6},\\
G_3 &= \max_{k\in\left[\left(\frac{m}{4}\right)^{\frac{1}{3}},\frac{m(\sqrt{m}-1)}{m-1}\right]} \frac{m/k}{\sqrt{m}/2} \leqslant \max_{k\in\left[\left(\frac{m}{4}\right)^{\frac{1}{3}},\frac{m(\sqrt{m}-1)}{m-1}\right]} 2\cdot\frac{\sqrt{m}}{k} \leqslant 2\cdot\frac{\sqrt{m}}{(m/4)^{1/3}} = 2^{5/3}\cdot m^{1/6},\\
G_4 &= \max_{k\geqslant\frac{m(\sqrt{m}-1)}{m-1}} \frac{m/k}{m/(k+m/k)} = \max_{k\geqslant\frac{m(\sqrt{m}-1)}{m-1}} 1+\frac{m}{k^2} = 1+\frac{m\cdot(m-1)^2}{m^2\cdot(\sqrt{m}-1)^2} = 1+\frac{(\sqrt{m}+1)^2}{m}\\
&\leqslant 5.
\end{aligned}$$

It is easy to check that $2^{8/3}\cdot m^{1/6} \leqslant 6.35\cdot m^{1/6}$ is the highest among all four factors, for all values of m. Hence, the upper and the lower bounds are tight by a factor of at most $6.35\cdot(m/4)^{1/6}$.

Regret, deterministic rules:

Upper bound: We show that the upper bound in this case is achieved by the rule f^* that selects the k alternatives with the highest plurality scores. Fix a profile $\vec{\sigma}$ and a set of alternatives $T \in \mathcal{A}_k \setminus \{f^*(\vec{\sigma})\}$. Let us calculate the worst-case regret due to T in the simplified regret formula from Lemma 6.1.

Recall that there are $\texttt{plu}(f^*(\vec{\sigma}),\vec{\sigma})$ votes i in which $\sigma_i(f^*(\vec{\sigma})) = 1$, and thus we cannot have $T \succ_i f^*(\vec{\sigma})$. Further, there are exactly $\texttt{plu}(T\setminus f^*(\vec{\sigma}),\vec{\sigma})$ votes i in which $T \succ_i f^*(\vec{\sigma})$ and $\sigma_i(T) = 1$. Hence, there are at most $n - \texttt{plu}(T\setminus f^*(\vec{\sigma}),\vec{\sigma}) - \texttt{plu}(f^*(\vec{\sigma}),\vec{\sigma})$ votes i in which $T \succ_i f^*(\vec{\sigma})$ and $\sigma_i(T) \geqslant 2$. Substituting these into the formula from Lemma 6.1, we get that the worst-case regret due to T is at most

$$\frac{1}{n}\cdot\left(\frac{\texttt{plu}(T\setminus f^*(\vec{\sigma}),\vec{\sigma})}{1} + \frac{n-\texttt{plu}(T\setminus f^*(\vec{\sigma}),\vec{\sigma})-\texttt{plu}(f^*(\vec{\sigma}),\vec{\sigma})}{2}\right) \tag{6.8}$$

$$= \frac{n+\texttt{plu}(T\setminus f^*(\vec{\sigma}),\vec{\sigma})-\texttt{plu}(f^*(\vec{\sigma}),\vec{\sigma})}{2n}. \tag{6.9}$$

Next, note that $\texttt{plu}(T\setminus f^*(\vec{\sigma}),\vec{\sigma}) \leqslant \texttt{plu}(T,\vec{\sigma}) \leqslant \texttt{plu}(f^*(\vec{\sigma}),\vec{\sigma})$, where the last equation follows due to the definition of f^*. Substituting this into Equation (6.9), we get that the regret is at most $1/2$, as desired.

For $k > m/2$, we can derive a better bound because we know $T \cap f^*(\vec{\sigma}) \neq \emptyset$. Note that $T \setminus f^*(\vec{\sigma}) \subseteq A \setminus f^*(\vec{\sigma})$. Because $f^*(\vec{\sigma})$ consists of the k alternatives with the highest plurality scores, and the plurality scores sum to n, we have

$$\texttt{plu}(f^*(\vec{\sigma}), \vec{\sigma}) \geqslant \frac{k}{m} \cdot n. \tag{6.10}$$

Similarly, $A \setminus f^*(\vec{\sigma})$ consists of the $m-k$ alternatives with the lowest plurality scores. Hence, we have $\texttt{plu}(A \setminus f^*(\vec{\sigma}), \vec{\sigma}) \leqslant (m-k) \cdot n/m$. Hence, we have

$$\texttt{plu}(T \setminus f^*(\vec{\sigma}), \vec{\sigma}) \leqslant \texttt{plu}(A \setminus f^*(\vec{\sigma}), \vec{\sigma}) \leqslant \frac{m-k}{m} \cdot n \tag{6.11}$$

Substituting Equations (6.10) and (6.11) into Equation (6.9), we get that the worst-case regret caused by T is at most

$$\frac{n + \frac{m-k}{m} \cdot n - \frac{k}{m} \cdot n}{2n} = 1 - \frac{k}{m}.$$

Since the choices of T and $\vec{\sigma}$ were arbitrary, we have that $\texttt{reg}(f^*) \leqslant 1 - k/m$, as required.

Lower bound: Next, we prove the matching lower bound. For $k \leqslant m/2$, fix a set $X \subseteq A$ of $2k$ alternatives. Construct a profile $\vec{\sigma}$ in which every alternative in X appears first in $n/(2k)$ votes. In this case, for any deterministic rule f with $|f(\vec{\sigma})| = k$, one can find a set $T \subseteq X \setminus f(\vec{\sigma})$ with $|T| = k$. Note that $T \cap f(\vec{\sigma}) = \emptyset$, and there are exactly $k \cdot n/(2k) = n/2$ votes in which $T \succ_i f(\vec{\sigma})$ and $\sigma_i(T) = 1$. Hence, due to Lemma 6.1, we have $\texttt{reg}(f) \geqslant \texttt{reg}(f, \vec{\sigma}) \geqslant (1/n) \cdot (n/2) = 1/2$, as required.

For $k > m/2$, we construct a profile $\vec{\sigma}$ in which every alternative appears first in exactly n/m votes. Once again, for any deterministic rule f with $|f(\vec{\sigma})| = k$, we can choose a set $T \supseteq A \setminus f(\vec{\sigma})$ with $|T| = k$. Note that $|T \setminus f(\vec{\sigma})| = m - k$. Hence, there are $(m-k) \cdot n/m$ votes in which $T \succ_i f(\vec{\sigma})$, and $\sigma_i(T) = 1$. Thus, due to Lemma 6.1, we have that $\texttt{reg}(f) \geqslant \texttt{reg}(f, \vec{\sigma}) \geqslant (1/n) \cdot (m-k) \cdot n/m = 1 - k/m$, as required.

Regret, randomized rules:

Upper bound: We explicitly construct the randomized rule f^* that provides the required upper bound. Fix a preference profile $\vec{\sigma}$. Without loss of generality, let us relabel the set of alternatives as $A = \{a_1, a_2, \ldots, a_m\}$ such that $plu(a_i, \vec{\sigma}) \geqslant plu(a_{i+1}, \vec{\sigma})$ for $i \in [m-1]$. Further, define $A_i \triangleq \{a_1, a_2, \ldots, a_i\}$ for $i \in [m]$.

Next, let t be the smallest positive integer such that

$$t \geqslant k + \texttt{plu}(a, \vec{\sigma}) \cdot \sum_{b \in A_t} \frac{1}{\texttt{plu}(b, \vec{\sigma})}, \quad \forall a \in A_t.$$

Next, define

$$p_a = \begin{cases} 1 - \frac{t-k}{\texttt{plu}(a,\vec{\sigma}) \cdot \sum_{b \in A_t} 1/\texttt{plu}(b,\vec{\sigma})} & \text{if } a \in A_t, \\ 0 & \text{otherwise.} \end{cases}$$

Observe that $0 \leqslant p_a \leqslant 1$ for each $a \in A$, and $\sum_{a \in A} p_a = k$. Hence, due to the bihierarchy extension [44] of the Birkhoff-von Neumann theorem [25, 197], there exists a distribution over $\mathcal{A}_k$ (which can be computed in polynomial time) under which the probability of each alternative $a \in A$ being selected is p_a. We let our rule f^* return this distribution.

Next, we bound $\mathtt{reg}(f^*)$. Fix $T \in \mathcal{A}_k$. Using Lemma 6.1, the worst-case regret due to T is 1 from each voter that ranks an alternative from $T \setminus f^*(\vec{\sigma})$ first, and at most $1/2$ from each voter that ranks an alternative from $A \setminus (f^*(\vec{\sigma}) \cup T)$ first. Hence, the expected regret is at most

$$\begin{aligned}
&\frac{1}{n} \cdot \mathbb{E}\left[\sum_{a \in T \setminus f^*(\vec{\sigma})} \mathtt{plu}(a, \vec{\sigma}) + \frac{1}{2} \sum_{a \in A \setminus (f^*(\vec{\sigma}) \cup T)} \mathtt{plu}(a, \vec{\sigma})\right] \\
&= \frac{1}{2n} \cdot \mathbb{E}\left[\sum_{a \in T \setminus f^*(\vec{\sigma})} \mathtt{plu}(a, \vec{\sigma})\right] + \frac{1}{2n} \cdot \mathbb{E}\left[\sum_{a \in A \setminus f^*(\vec{\sigma})} \mathtt{plu}(a, \vec{\sigma})\right] \\
&= \frac{1}{2n} \cdot \sum_{a \in T} (1 - p_a) \cdot \mathtt{plu}(a, \vec{\sigma}) + \frac{1}{2} - \frac{1}{2n} \cdot \sum_{a \in A_t} p_a \cdot \mathtt{plu}(a, \vec{\sigma}) \\
&\leqslant \frac{k \cdot (t-k)}{2n \cdot \sum_{a \in A_t} 1/\mathtt{plu}(a, \vec{\sigma})} + \frac{1}{2} - \frac{1}{2n} \cdot \sum_{a \in A_t} \mathtt{plu}(a, \vec{\sigma}) + \frac{t \cdot (t-k)}{2n \cdot \sum_{a \in A_t} 1/\mathtt{plu}(a, \vec{\sigma})} \\
&= \frac{1}{2} + \frac{t^2 - k^2}{2n \cdot \sum_{a \in A_t} 1/\mathtt{plu}(a, \vec{\sigma})} - \frac{1}{2n} \cdot \sum_{a \in A_t} \mathtt{plu}(a, \vec{\sigma}) \\
&\leqslant \frac{1}{2} - \frac{k^2}{2n} \cdot \sum_{a \in A_t} \mathtt{plu}(a, \vec{\sigma}) \leqslant \frac{1}{2} - \frac{k^2}{2n \cdot t \cdot (m/n)} \leqslant \frac{1}{2} \cdot \left(1 - \frac{k^2}{m^2}\right),
\end{aligned}$$

as desired. The first two equalities follow from the definitions of f^*, T, A_t, and p_a. The first inequality follows because our definitions guarantee

$$(1 - p_a) \cdot \mathtt{plu}(a, \vec{\sigma}) \leqslant \frac{t - k}{\sum_{b \in A_t} 1/\mathtt{plu}(b, \vec{\sigma})}$$

for every $a \in A$. The second inequality follows by the power mean inequality, the third inequality follows because the t alternatives in A_t have the highest plurality scores (hence, $\sum_{a \in A_t} \mathtt{plu}(a, \vec{\sigma}) \geqslant t \cdot n/m$), and the final inequality follows because $t \leqslant m$.

Lower bound: This proof is very similar to the proof of the deterministic case.

For $k \leqslant m/2$, fix a set of alternatives $X \subseteq A$ such that $|X| = 2k$. Construct a profile $\vec{\sigma}$ in which every alternative in X appears first in $n/(2k)$ votes. Now, let us consider a randomized rule f that returns a distribution over $\mathcal{A}_k$. Let T denote the set of k alternatives from X with the least probability of being picked in $\mathcal{S}$. Because $\sum_{a \in X} \Pr[a \in f(\vec{\sigma})] \leqslant k$, we have $\sum_{a \in T} \Pr[a \in f(\vec{\sigma})] \leqslant k/2$. Hence, $\sum_{a \in T} \Pr[a \notin f(\vec{\sigma})] \geqslant k/2$. Now, from Lemma 6.1, the expected regret of f due to T is at least

$$\frac{1}{n} \cdot \mathbb{E}\left[\sum_{a \in T} \mathbb{I}[a \notin f(\vec{\sigma})] \cdot \frac{n}{2k}\right] = \frac{1}{2k} \cdot \sum_{a \in T} \Pr[a \notin f(\vec{\sigma})] \geqslant \frac{1}{4},$$

as required.

Similarly, for $k > m/2$, once again construct a profile $\vec{\sigma}$ in which every alternative appears first in n/m votes. Consider a randomized rule f that returns a distribution over $\mathcal{A}_k$. Let T be the set of k alternatives with the least probability of being picked in $f(\vec{\sigma})$. Because $\sum_{a \in A} \Pr[a \in f(\vec{\sigma})] \leqslant k$, we have $\sum_{a \in T} \Pr[a \in f(\vec{\sigma})] \leqslant k \cdot k/m$. Hence, $\sum_{a \in T} \Pr[a \notin f(\vec{\sigma})] \geqslant k \cdot (1 - k/m)$. Once again, from Lemma 6.1, the expected regret of f due to T is at least

$$\frac{1}{n} \cdot \mathbb{E}\left[\sum_{a \in T} \mathbb{I}[a \notin f(\vec{\sigma})] \cdot \frac{n}{m}\right] = \frac{1}{m} \cdot \sum_{a \in T} \Pr[a \notin f(\vec{\sigma})] \geqslant \frac{k}{m} \cdot \left(1 - \frac{k}{m}\right),$$

as required.

Gap between upper and lower bounds: Note that for $k \leqslant m/2$, the ratio between the upper and the lower bounds is

$$\frac{(1/2) \cdot (1 - k^2/m^2)}{1/4} = 2 \cdot \left(1 - \frac{k^2}{m^2}\right) \leqslant 2.$$

For $k > m/2$, the ratio between the upper and the lower bounds is

$$\frac{(1/2) \cdot (1 - k^2/m^2)}{(1/2) \cdot (k/m) \cdot (1 - k/m)}.$$

It can be checked easily that this is a decreasing function of k. Hence, the maximum ratio is achieved at $k = m/2$, and is equal to 2.

Thus, in both cases, the upper and the lower bounds in this case are tight up to a constant factor of 2.

Running time: Note that rules from our upper bounds only require calculating the plurality scores and finding a decomposition according to (the bihierarchy extension of) the Birkhoff-von Neumann theorem, both of which can be accomplished in polynomial time. ■

6.4 Empirical Comparisons

In Section 6.3 we provided analytical results for both deterministic and randomized rules. In our view, randomized rules are especially practicable when the output distribution is sampled multiple times, or when the voters are well-informed, or when the voters are indifferent about the outcome (e.g., they are software agents). Moreover, we believe that the results for randomized rules are of substantial theoretical interest. But our work is partly driven by its direct applications in RoboVote (see Section 6.1.1), which does not satisfy the above conditions. This leads us to use deterministic voting rules, which is what we focus on hereinafter.

Let $f^*_{\texttt{dist}}$ and $f^*_{\texttt{reg}}$ be the *deterministic* rules that minimize the worst-case distortion and regret, respectively, on every given preference profile. The deterministic results of

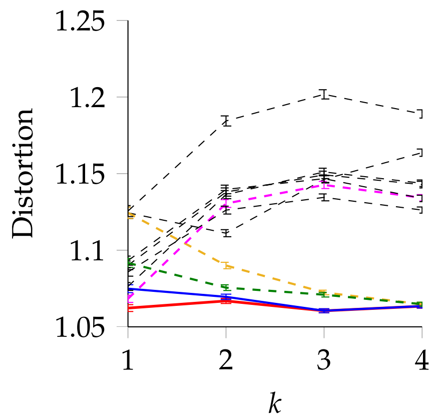

(a) Avg. Distortion

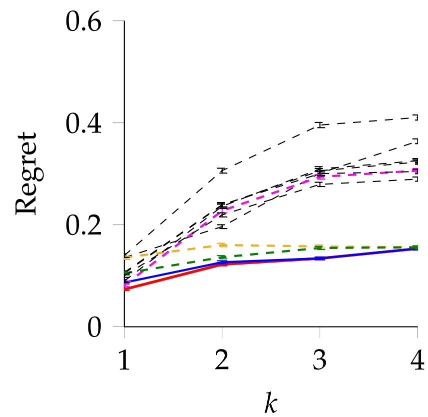

(b) Avg. Regret

Figure 6.2: Uniformly random utility profiles.

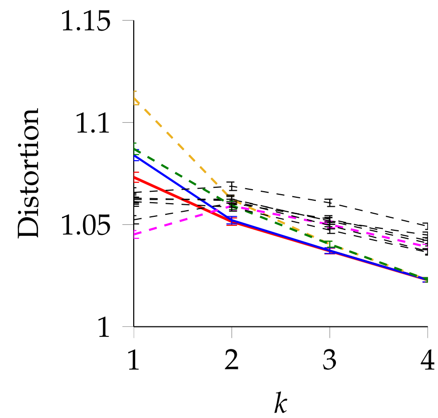

(a) Avg. Distortion

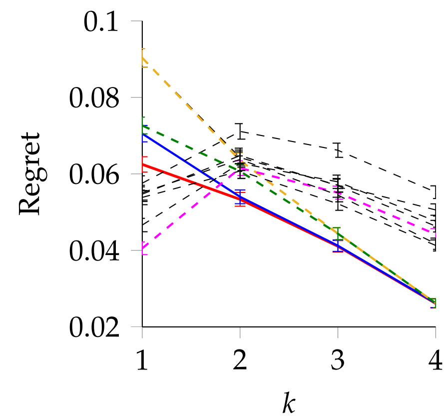

(b) Avg. Regret

Figure 6.3: Utility profiles from the Jester dataset.

Section 6.3 establish upper and lower bounds on their *worst-case* distortion/regret. In this section, we evaluate their *average-case* performance on simulated as well as real data, and compare them against nine well-known voting rules: plurality, approval voting, Borda count, STV, Kemeny's rule, the maximin rule, Copeland's rule, Bucklin's rule, and Tideman's rule.[4]

We perform three experiments: (i) choosing a utility profile uniformly at random from the simplex of all utility profiles, (ii) drawing a real-world utility profile from the Jester datasets [103], and (iii) drawing a real-world preference profile from the PrefLib datasets [141], and choosing a consistent utility profile uniformly at random. For each

[4]For the score-based rules, the k-subset is selected by picking the top k alternatives based on their scores.

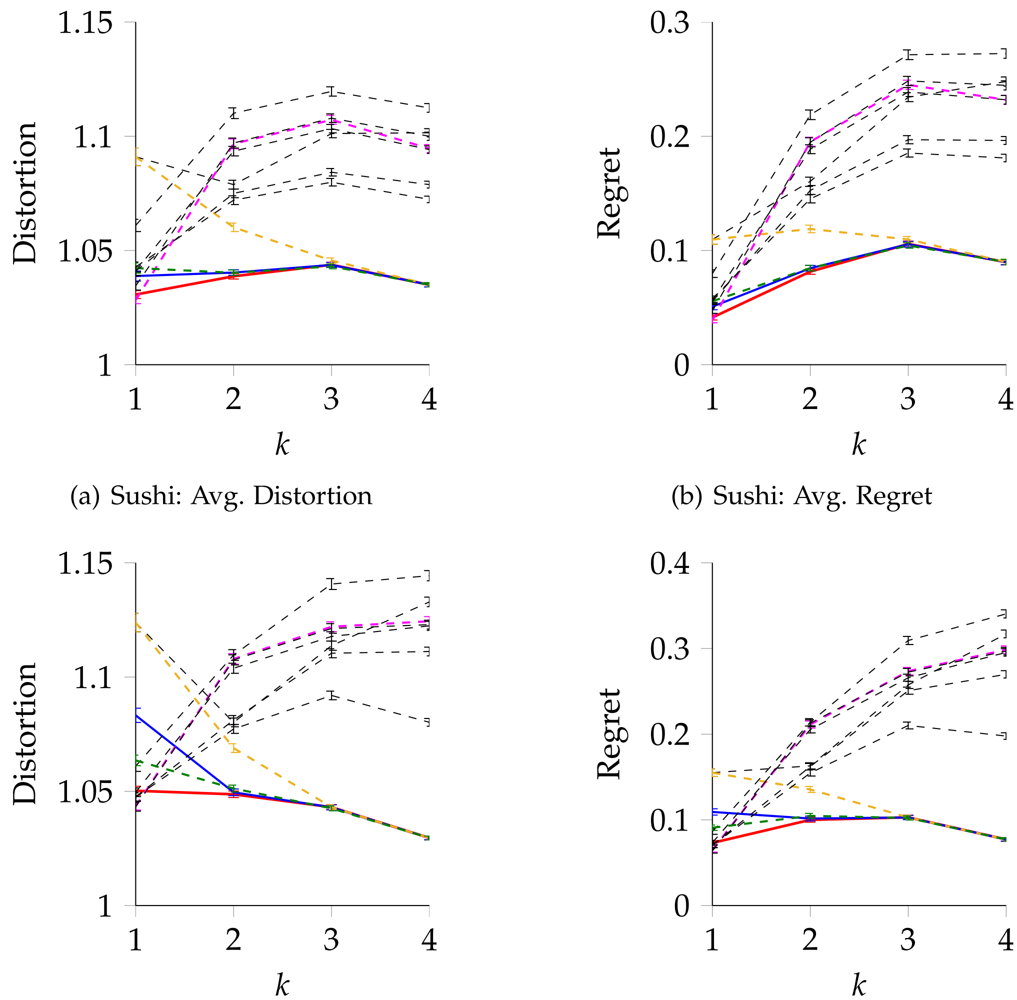

(a) Sushi: Avg. Distortion

(b) Sushi: Avg. Regret

(c) T Shirt: Avg. Distortion

(d) T Shirt: Avg. Regret

Figure 6.4: Preference profiles from Sushi and T-Shirt datasets, uniformly random consistent utility profiles.

experiment, we have 8 voters and 10 alternatives, and test for $k \in [4]$.[5] For each setting, we perform 10 000 random simulations, and measure both distortion and regret for the *actual utility profile*, as opposed to the worst-case utility profile. The figures show the average performance with 95% confidence intervals.

In all of our simulations, we observed that three of the classical voting rules stand out: Borda count performs well for choosing a single alternative (but not for choosing larger subsets) whereas plurality and STV perform well for choosing larger subsets (but not for choosing a single alternative). Hence, all of our graphs specifically distinguish these three rules in addition to $f^*_{\texttt{dist}}$ and $f^*_{\texttt{reg}}$.

[5]In RoboVote, we expect typical instances to have few voters and alternatives.

Figure 6.2 shows the results for the first experiment where we choose the utility profile uniformly at random. Figure 6.3 shows the results for the second experiment where real-world utility profiles are drawn from one of the Jester datasets, in which more than 50 000 voters rated 150 jokes on a real-valued scale; the results from the other Jester dataset are almost identical. Finally, Figure 6.4 shows the results for the third experiment where real-world preference profiles are drawn from the Sushi dataset (5 000 voters ranking 100 different kinds of sushi) and the T-Shirt dataset (30 voters ranking 11 T-shirt designs) from PrefLib. Experiments on other datasets from PrefLib (AGH Course Selection, Netflix, Skate, and Web Search) yielded similar results.

Right off the bat, one can observe that the average-case distortion and regret values are much lower than their worst-case counterparts. For example, average regret is generally lower than 0.1 — compare with the tight worst-case deterministic bound of $1/2$ for $k \leqslant m/2$.

Much to our surprise, in all of our experiments, $f^*_{\texttt{reg}}$ outperforms $f^*_{\texttt{dist}}$ in terms of both average-case distortion (multiplicative loss) *and* regret (additive loss). While both measures of loss have been studied extensively in the literature, we are not aware of any previous work that compares the two approaches. At least in our social choice domain, the regret-based approach is clearly better on average.

Moreover, in all cases but one ($k = 1$ in the Jester experiment), $f^*_{\texttt{reg}}$ also outperforms all the classical voting rules under consideration. We therefore conclude that, on random as well as on real-world instances, $f^*_{\texttt{reg}}$ provides superior performance in terms of social welfare maximization.

6.5 Computation and Implementation

In this section, we analyze and compare the two deterministic optimal rules — $f^*_{\texttt{dist}}$ and $f^*_{\texttt{reg}}$ — from a computational viewpoint. Selecting optimal subsets turns out to be challenging, as both rules are $\mathcal{NP}$-hard to compute.

Theorem 6.2. *Given a preference profile $\vec{\sigma}$ and an integer k, computing a k-subset of alternatives that has the minimum distortion or the minimum regret on $\vec{\sigma}$ is $\mathcal{NP}$-hard.*

Proof. We present a polynomial-time reduction from the minimum dominating set problem, which is known to be NP-hard (in fact, APX-hard) even in 3-regular graphs (e.g., see [161]). A set of nodes S is called a dominating set in a graph $G = (V, E)$ if any node in $V \setminus S$ is adjacent to a node in S, and is called a minimum dominating set if it is a dominating set of the minimum possible size.

Let $V = \{v_1, \ldots, v_n\}$ (thus, $n = |V|$). Observe that 3-regularity of G implies $|E| = 3n/2$. Let c and d be positive integers such that c is a multiple of $12n$, $c \geqslant 96n^2 H_d$ (where H_d is the d^{th} harmonic number), and $d = 3nc + 3c + \sum_{i=2}^{3n/2} \left(5c - \frac{c \cdot i}{3n}\right)$. Clearly, there exist such values of c and d satisfying $c = \mathcal{O}(n^2 \ln n)$ and $d = \mathcal{O}(n^3 \ln n)$, and they can be computed in polynomial time.

Now, set $k = 1 + \lfloor 3n/4 \rfloor$, and construct a profile $\vec{\sigma}$ as follows. The set of alternatives is the union of three sets:

- $A = \{a_1, \ldots, a_n\}$ (each a_i corresponds to the node $v_i \in V$);
- $B = \{b_1, \ldots, b_{3n/2}\}$ (each b_i corresponds to an edge in E);
- $D = \{f_1, \ldots, f_d\}$.

The set of voters N is the union of the following sets:

- N_e, which consists of $2c$ "edge voters" $\left\{n_e^j\right\}_{j \in \{0,1,\ldots,2c-1\}}$ for each edge $e \in E$;
- N_1, which consists of $3c$ voters;
- N_i consisting of $\frac{14c}{3} - \frac{c \cdot i}{12n}$ voters, for each $i = 2, \ldots, 3n/2$.

The reader may check that the total number of voters is exactly d. Next, the votes in $\vec{\sigma}$ are as follows:

- For each edge $e = (v_{i_1}, v_{i_2}) \in E$ with $i_1 < i_2$, voter n_e^j ranks alternatives a_{i_1} and a_{i_2} at positions 2 and 3, respectively, when j is even, and at positions 3 and 2, respectively, when j is odd. Alternatives in $A \setminus \{a_{i_1}, a_{i_2}\} \cup B$ are ranked at the bottom, in arbitrary order. Positions 1 and 4 through $d + 2$ are reserved for alternatives in D (see below for the way the alternatives in D are placed in these positions).
- Voters in N_1 rank alternative b_1 in the first position, and alternatives in $A \cup B \setminus \{b_1\}$ in the last positions, in arbitrary order. Positions 2 through $d + 1$ are reserved for alternatives in D.
- For $i = 2, \ldots, 3n/2$, voters in N_i rank alternative b_i in the second position, and alternatives in $A \cup B \setminus \{b_i\}$ in the last positions, in arbitrary order. Positions 1 and 3 through $d + 1$ are reserved for alternatives in D.
- Alternatives in D are shuffled in a cyclic fashion in the votes within the positions reserved for them. More specifically, fix an order among the votes. Then, for $i, j \in [d]$, let $t = 1 + (i + j - 2 \bmod d)$. Alternative f_i appears in the t^{th} reserved position of the j^{th} vote.

Given this construction, the next lemma establishes a strong relation between the k-sized set of alternatives with the minimum regret or distortion in $\vec{\sigma}$, and the minimum dominating set in graph G. Theorem 6.2 then follows immediately because our reduction is polynomial time.

Lemma 6.2. *A k-sized set of alternatives has the minimum regret or the minimum distortion on $\vec{\sigma}$ if and only if it consists of the alternatives in A corresponding to the nodes of a minimum dominating set S^* of G and the alternatives $b_1, b_2, \ldots, b_{k-|S^*|}$.*

Proof. Let K^* denote the k-sized set of alternatives that consists of the alternatives in A corresponding to the nodes of a minimum dominating set S^* of G, and the alternatives $b_1, b_2, \ldots, b_{k-|S^*|}$. We prove an upper and a lower bound on the regret of K^* on $\vec{\sigma}$, which establishes that it is the unique k-sized set of alternatives with the minimum regret. We later provide the argument that shows it is also the unique k-sized set of alternatives with the minimum distortion.

Upper bound: Let K' be a k-sized set of alternatives that is disjoint from K^*.[6]. We now show that the regret of K^* on $\vec{\sigma}$ due to K' is at most

$$\texttt{reg}(K^*, \vec{\sigma}) \leqslant T \triangleq \sum_{i=k-|S^*|+1}^{2k-|S^*|} \left(\frac{7c}{3} - \frac{c \cdot i}{24n} + 1 \right). \tag{6.12}$$

Since T is independent of K' itself, it would follow that the worst-case regret of K^* is also upper bounded by T.

Consider an alternative $a_i \in A \setminus K^*$. Recall that this corresponds to the node $v_i \in V$. Let v_ℓ be a node in S^* that is adjacent to v_i. From Equation 6.1, the contribution of alternative a_i to the regret is

- 0 in the c edge voters n_e^j corresponding to edge $e = (v_i, v_\ell)$ that have alternative a_i ranked third;
- at most $1/3$ in each of the remaining $2c$ edge voters that have alternative a_i ranked third;
- at most $1/2$ in each of the $3c$ edge voters that have alternative a_i ranked second;
- smaller than $1/d$ in any other voter (because it is ranked below all alternatives in D in these voters).

Thus, the total contribution of a_i to the regret is at most $2c/3 + 3c/2 + 1 = 13c/6 + 1$.

Because of the cyclic shuffling, the contribution of an alternative in D to the regret is at most H_d.

For $i \geqslant k + 1 - |S^*|$, the contribution of alternative b_i to the regret is at most $1/2$ from the voters in N_i, and at most $1/d$ in any other voter. In total, this contribution is in $[7c/3 - \frac{c \cdot i}{24n}, 7c/3 - \frac{c \cdot i}{24n} + 1]$, and therefore is higher than the contribution of any alternative in $(A \cup D) \setminus K^*$.

Hence, the k alternatives that contribute the highest regret are $b_{k+1-|S^*|}, \ldots, b_{2k-|S^*|}$, and their total contribution to the regret is at most $\sum_{i=k+1-|S^*|}^{2k-|S^*|} \left(\frac{7c}{3} - \frac{c \cdot i}{24n} + 1 \right)$ as desired.

Lower bound: Next, let K be a k-sized set of alternatives that does not follow the characterization of the lemma (i.e., does not consist of exactly the alternatives in A corresponding to the nodes of a minimum dominating set S^* of G and the alternatives $b_1, \ldots, b_{k-|S^*|}$). We show that $\texttt{reg}(K, \vec{\sigma}) > T$. This would establish that K^* is the unique k-sized alternative with the minimum regret.

Let S be the set of nodes of G corresponding to the alternatives in $K \cap A$. Let $\mathrm{I}(S)$ denote an independent set of nodes of G (i.e., no two nodes in $\mathrm{I}(S)$ are connected to each other) such that no node in $\mathrm{I}(S)$ is adjacent to any node in S. Construct a k-sized set K' that consists of the alternatives in A corresponding to the nodes of $\mathrm{I}(S)$ (if any), and the $k - |\mathrm{I}(S)|$ alternatives from $B \setminus K$ with the smallest indices. We show that the regret of K on $\vec{\sigma}$ due to K' is more than T.

[6]If K' shares alternatives with K^*, such alternatives do not cause any regret in the simplified regret formula from Equation (6.1). Hence, to compute the worst-case regret, it is sufficient to focus on sets K' that are disjoint from K^*.

We do so by considering the contribution of the alternatives of K' to the regret only from the votes in which they appear among the first three positions. For now, let us ignore the alternatives in D that appear in K; we will later the regret calculated to account for such alternatives.

Let $X = \mathbb{I}[(B \setminus K) \cap \{b_1, \ldots, b_{k-|S|-|K\cap D|}\} \neq \emptyset]$. Consider an alternative a_i corresponding to the node $v_i \in \mathrm{I}(S)$. Because the nodes in $\mathrm{I}(S)$ are mutually non-adjacent and non-adjacent to any node in S, $a_i \in K'$ is ranked above all alternatives in $K \setminus D$ in the edge voters n_e^j for $j = 0, 1, \ldots, 2c-1$ and each edge e adjacent to node v_i. Ignoring the alternatives in $K \cap D$ (that may be ranked in the first position of these voters), the contribution of a_i to the regret is $1/2$ from each of the $3c$ edge voters that rank a_i in the second position, and $1/3$ from each of the $3c$ edge voters that rank a_i in the third position. Thus, the total contribution of the alternatives that correspond to the nodes in $\mathrm{I}(S)$ is $\frac{5c}{2}$.

Next, consider the $k - |\mathrm{I}(S)|$ alternatives from $B \setminus K$ that have the smallest indices. If K contains all of the alternatives $b_1, \ldots, b_{k-|S|-|K\cap D|}$, then K' contains the alternatives b_i for $i = k - |S| - |K \cap D| + 1, \ldots, 2k - |S| - |K \cap D| - |\mathrm{I}(S)|$. Again, ignoring the alternatives in $K \cap D$ (that may be ranked in the first position in the votes of the voters in N_i), the contribution of b_i to the regret of K is $1/2$ from each of the $\frac{14c}{3} - \frac{c \cdot i}{12n}$ voters in N_i that rank b_i in the second position.

If $X = 1$, we know that K does not use the alternatives of B with the smallest indices, and the contribution of the alternatives of B to the regret increases by at least $\frac{c \cdot X}{24n}$.

Finally, let us consider the fact that K may include some alternatives from D. Each such alternative may be ranked first by a voter, for which we have incorrectly added a regret of at most $1/2$. Hence, the actual regret may be $(1/2) \cdot |K \cap D|$ lower than our calculated regret.

Combining the entire analysis, and observing that the regret due to K' is a lower bound on the worst-case regret of K on σ, we get that

$$
\begin{aligned}
&\mathtt{reg}(K, \vec{\sigma}) \\
&\quad \geqslant \frac{5c}{2} \cdot |\mathrm{I}(S)| + \sum_{i=k-|S|-|K\cap D|+1}^{2k-|S|-|K\cap D|-|\mathrm{I}(S)|} \left(\frac{7c}{3} - \frac{c \cdot i}{24n} \right) + \frac{c \cdot X}{24n} - \frac{1}{2} \cdot |K \cap D| \\
&\quad \geqslant \frac{c}{6} \cdot |\mathrm{I}(S)| + \sum_{i=k-|S|-|K\cap D|+1}^{2k-|S|-|K\cap D|} \left(\frac{7c}{3} - \frac{c \cdot i}{24n} \right) + \frac{c \cdot X}{24n} - \frac{1}{2} \cdot |K \cap D| \\
&\quad = \sum_{i=k-|S^*|+1}^{2k-|S^*|} \left(\frac{7c}{3} - \frac{c \cdot i}{24n} + 1 \right) + \frac{c}{6} \cdot |\mathrm{I}(S)| - \frac{c \cdot k}{24n} \cdot (|S^*| - |S|) \\
&\qquad\qquad + \left(\frac{c \cdot k}{24n} - \frac{1}{2} \right) \cdot |K \cap D| - k + \frac{c \cdot X}{24n} \\
&\quad \geqslant \mathtt{reg}(K^*, \vec{\sigma}) + \frac{c}{6} \cdot |\mathrm{I}(S)| - \frac{c \cdot k}{24n} \cdot (|S^*| - |S|) + \left(\frac{c \cdot k}{24n} - \frac{1}{2} \right) \cdot |K \cap D| - k + \frac{c \cdot X}{24n}
\end{aligned}
\tag{6.13}
$$

Finally, we analyze the quantities on the RHS of Equation (6.13) to derive that

$$\texttt{reg}(K,\vec{\sigma}) - \texttt{reg}(K^*,\vec{\sigma}) > \frac{7k}{9} \cdot |K \cap D|, \tag{6.14}$$

which implies that $\texttt{reg}(K,\vec{\sigma}) > \texttt{reg}(K^*,\vec{\sigma})$, as required.

We now perform a case-by-case analysis to establish Equation (6.14). First, suppose S is a dominating set of G.

- If S is a minimum dominating set of G (i.e., $|S| = |S^*|$), we must have $X = 1$ or $|K \cap D| > 0$.
  - If $X = 1$, Equation (6.13) yields

    $$\texttt{reg}(K,\vec{\sigma}) - \texttt{reg}(K^*,\vec{\sigma}) \geqslant \left(\frac{c \cdot k}{24n} - \frac{1}{2}\right) \cdot |K \cap D| - k + \frac{c}{24n},$$

    which in turn implies Equation (6.14) because the definition of c and the fact that $k \geqslant 1$ imply $\frac{c}{24n} > k$ and $(c \cdot k)/(24n) - (1/2) > (7k)/9$.
  - If $X = 0$ and $|K \cap D| > 0$, Equation (6.13) yields

    $$\begin{aligned}\texttt{reg}(K,\vec{\sigma}) - \texttt{reg}(K^*,\vec{\sigma}) &\geqslant \left(\frac{c \cdot k}{24n} - \frac{1}{2}\right) \cdot |K \cap D| - k \\ &\geqslant \left(\frac{c}{24n} - \frac{3}{2}\right) \cdot k \cdot |K \cap D| > \frac{7k}{9} \cdot |K \cap D|,\end{aligned}$$

    where the last inequality follows from the definition of c.
- If $|S| > |S^*|$, Equation (6.13) yields

  $$\texttt{reg}(K,\vec{\sigma}) - \texttt{reg}(K^*,\vec{\sigma}) \geqslant \left(\frac{c}{24n} - 1\right) \cdot k + \left(\frac{c \cdot k}{24n} - \frac{1}{2}\right) \cdot |K \cap D| > \frac{7k}{9} \cdot |K \cap D|,$$

  where the last inequality follows from the definition of c.

Next, suppose S is not a dominating set of G. Here, we distinguish between two cases.

- If $|S| \geqslant |S^*|$, using $|\mathrm{I}(S)| \geqslant 1$, Equation (6.13) yields

  $$\texttt{reg}(K,\vec{\sigma}) - \texttt{reg}(K^*,\vec{\sigma}) \geqslant \frac{c}{6} + \left(\frac{c \cdot k}{24n} - 1\right) \cdot |K \cap D| - k.$$

  Equation (6.14) now follows because the definition of c implies $(c/6) > k$ and $(c \cdot k)/(24n) - 1 > (7k)/9$.
- If $|S| < |S^*|$, then we claim that $|\mathrm{I}(S)| \geqslant (|S^*| - |S|)/4$. Indeed, because the minimum dominating set of G has size $|S^*|$, the set of nodes S must leave $|S^*| - |S|$ nodes of G undominated. Because G is 3-regular, any independent set $\mathrm{I}(S)$ among these nodes should have size at least $(|S^*| - |S|)/4$. Hence,

  $$\texttt{reg}(K,\vec{\sigma}) - \texttt{reg}(K^*,\vec{\sigma}) \geqslant \left(\frac{c}{24} - \frac{c \cdot k}{24n}\right) \cdot (|S^*| - |S|) + \left(\frac{c \cdot k}{6n} - 1\right) \cdot |K \cap D| - k$$

$$> \frac{c}{24} - \frac{c \cdot k}{24n} - k + \left(\frac{c}{24n} - \frac{3}{2}\right) \cdot k \cdot |K \cap D|,$$

which implies Equation (6.14) because the definition of c implies $(c/24) - (c \cdot k)/(24n) > k$ and $c/(24n) - (3/2) > (7/9)$.

This completes the proof of Equation (6.14) in all cases, and hence, concludes the proof of the regret part of the lemma.

Let us now prove the distortion part of the lemma. Again, consider a k-sized set K that does not follow the characterization of the lemma (i.e., does not consist of exactly alternatives in A corresponding to the nodes of a minimum dominating set S^* of G and the alternatives $b_1, \ldots, b_{k-|S^*|}$).

If $K \cap D = \emptyset$, then we have $\texttt{plu}(K, \vec{\sigma}) \leqslant \texttt{plu}(K^*, \vec{\sigma})$. Hence,

$$\texttt{dist}(K, \vec{\sigma}) = 1 + m \cdot \frac{\texttt{reg}(K, \vec{\sigma})}{\texttt{plu}(K, \vec{\sigma})} > 1 + m \cdot \frac{\texttt{reg}(K^*, \vec{\sigma})}{3c} = \texttt{dist}(K^*),$$

where the second transition uses Equation (6.14).

If $K \cap D \neq \emptyset$, we have $\texttt{plu}(K, \vec{\sigma}) \leqslant 3c + |K \cap D|$, which, together with Equation (6.14), implies

$$\begin{aligned}\texttt{dist}(K, \vec{\sigma}) = 1 + m \cdot \frac{\texttt{reg}(K, \vec{\sigma})}{\texttt{plu}(K, \vec{\sigma})} > 1 + m \cdot \frac{\texttt{reg}(K^*, \vec{\sigma}) + \frac{7k}{9} \cdot |K \cap D|}{3c + |K \cap D|} &> 1 + m \cdot \frac{\texttt{reg}(K^*)}{3c} \\ &= \texttt{dist}(K^*, \vec{\sigma}),\end{aligned}$$

where the final transition holds due to Equation (6.12), which implies $7k/9 > \texttt{reg}(K^*, \vec{\sigma})/(3c)$. ■ (Proof of Lemma 6.2)

This concludes the entire proof. ■ (Proof of Theorem 6.2)

Given that $f^*_{\texttt{reg}}$ outperforms $f^*_{\texttt{dist}}$ in the experiments of Section 6.4, and that both rules are computationally hard, $f^*_{\texttt{reg}}$ stands out as the clear choice for implementation in our website RoboVote. We therefore devoted our efforts to developing a scalable implementation for $f^*_{\texttt{reg}}$.

First, let us note the simplified formula for $f^*_{\texttt{reg}}$ that follows from Lemma 6.1:

$$f^*_{\texttt{reg},k}(\vec{\sigma}) = \operatorname*{arg\,min}_{T \in \mathcal{A}_k} \max_{S \in \mathcal{A}_k} \sum_{i=1}^{n} \frac{\mathbb{I}[S \succ_{\sigma_i} T]}{\sigma_i(S)}. \tag{6.15}$$

To better understand this formula, we consider the special case of $k = 1$. In this case,

$$f^*_{\texttt{reg}}(\vec{\sigma}) \in \operatorname*{arg\,min}_{a \in A} \max_{b \in A} \sum_{i=1}^{n} \frac{\mathbb{I}[b \succ_{\sigma_i} a]}{\sigma_i(b)}.$$

Note that this voting rule is very similar to the classical maximin rule: replacing $\sigma_i(b)$ with 1 in the denominator would yield the maximin rule. Thus, in some sense, this is

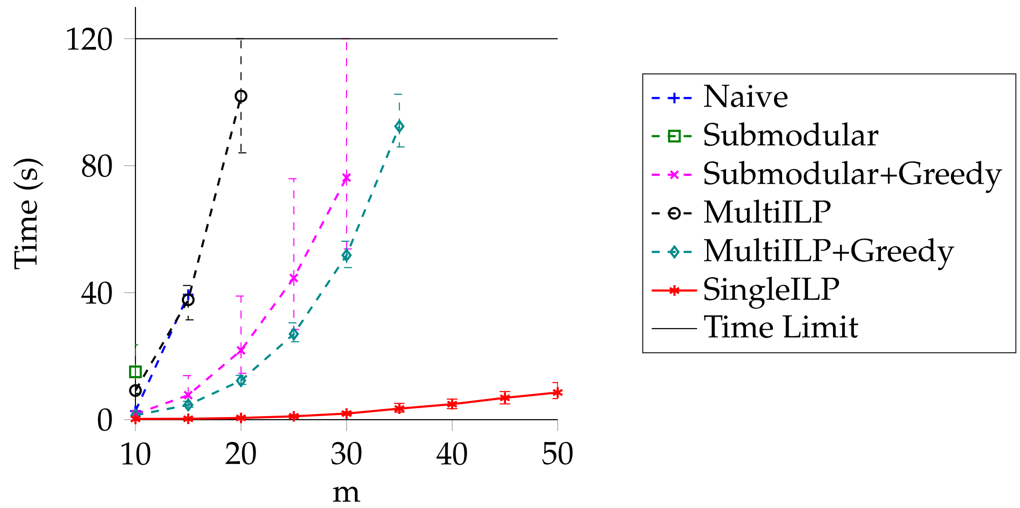

Figure 6.5: Running times of six approaches to computing $f^*_{\texttt{reg}}$.

a smooth version of the maximin rule, where the "victory" of b over a in voter i's vote is weighted by the strength of b in this vote (measured by $1/\sigma_i(b)$). In our view, this intuitive structure makes $f^*_{\texttt{reg}}$ even more compelling.

We now briefly describe six approaches we have developed for computing $f^*_{\texttt{reg}}$:

1. Naïve: This uses Equation (6.15), and requires $\Omega(n \cdot \binom{m}{k}^2)$ operations, which is prohibitive even for small m.
2. Submodular: The regret for set S in choosing set T, i.e., $\sum_{i \in [n]: S \succ_{\sigma_i} T} 1/\sigma_i(S)$, is submodular in S. Hence, for each $T \in \mathcal{A}_k$ we can optimize over $S \in \mathcal{A}_k$ using any algorithm for the submodular maximization subject to cardinality constraint (SMCC) problem. We use the SFO toolbox for Matlab [125].
3. Submodular+Greedy: This improves the previous approach by first computing a $1 - 1/e$ greedy approximation to the SMCC instance for set T, and pruning T if this is already greater than the best regret found so far.
4. MultiILP: Instead of using SMCC, for each $T \in \mathcal{A}_k$ we optimize over $S \in \mathcal{A}_k$ by solving an integer linear program (ILP) with roughly $n \cdot m$ variables and $n \cdot m^2$ constraints. Note that $\binom{m}{k}$ such ILPs need to be solved.
5. MultiILP+Greedy: This improves the MultiILP approach by using a greedy pruning procedure as before.
6. SingleILP: This approach solves a single but huge ILP with $\binom{m}{k}$ additional constraints.

Figure 6.5 shows the average running times of these approaches (and 95% confidence intervals) over 10 000 instances with $n = 15$, $k = 3$, and m varying from 10 to 50.[7] The experiments were performed on a single machine with quad-core 2.9 GHz CPU and 32 GB RAM. A time limit of 2 minutes was set because a running time greater than this

[7]The running time scales linearly in n, and increases with $\binom{m}{k}$.

would not be helpful for our website, where the results need to be delivered quickly to the users. While the greedy pruning procedure does help reduce the running time of both the Submodular and MultiILP approaches, SingleILP still computes $f^*_{\mathtt{reg}}$ much faster than any other approach, solving instances with 50 alternatives in less than 10 seconds. We have therefore implemented SingleILP on RoboVote.

6.6 Related Work

In addition to the work of Procaccia and Rosenschein [171] and Boutilier et al. [35], several other papers employ the notion of distortion to quantify how close one can get to maximizing utilitarian social welfare when only ordinal preferences are available [3, 4, 47]. In particular, Anshelevich et al. [4] study the same setting as Boutilier et al. [35], but in addition assume the preferences of voters are consistent with distances in a metric space. We refer the reader to the paper by Boutilier et al. [35, Section 1.2] for a thorough discussion of work (in philosophy, economics, and social choice theory) related to implicit utilitarian voting more broadly.

There is quite a bit of work in computational social choice on voting rules that select subsets of alternatives. Typically it is assumed that ordinal preferences are translated into a position-based score for each alternative (in contrast to our work). Just to give a few examples, under the Chamberlin-Courant method, each voter assigns a score to a set equal to the highest score of any alternative in the set, and the (computationally hard) objective is to choose a subset of size k that maximizes the sum of scores [54, 178]. Skowron et al. [191] generalize the way in which the score of a voter for a subset of alternatives is computed. Aziz et al. [12] propose selecting a subset of alternatives in order to satisfy a fairness axiom they term *justified representation*, and study whether common voting rules satisfy this axiom. The *budgeted social choice* framework of Lu and Boutilier [138] is more general in that the number of alternatives to be selected is not fixed; rather, each alternative has a cost that must be paid to add it to the selection.

6.7 Discussion

We find it exciting that new theoretical questions in computational social choice are driven by concrete real-world applications. And while research in the field is often motivated by potential applications to multiagent systems, we focus on helping *people* — not software agents — make joint decisions.

We also remark that we consider the empirical dominance of $f^*_{\mathtt{reg}}$, in terms of both regret and (surprisingly) distortion, to be especially significant. It would be interesting to understand, on a theoretical level, why this happens. A promising starting point is to derive analytical bounds on the average-case distortion of $f^*_{\mathtt{dist}}$ and $f^*_{\mathtt{reg}}$ under uniformly random utility profiles.

Chapter 7

Robust Voting Rules

7.1 Introduction

In the previous chapter, we studied the classic paradigm of social choice theory in which subjective preferences of individuals are aggregated towards a collective outcome. A different — in a sense competing — paradigm views voting rules as *estimators*.

From this viewpoint, there exists an underlying objective ground truth ranking that compares the alternatives, and the input votes are simply noisy estimates of this true ranking. Under this paradigm, one voting rule is better than another if it is more likely to output the true ranking. The prevalent approach in this paradigm is the *maximum likelihood estimator* (MLE) approach, which assumes a statistical model for the generation of noisy votes given the true ranking (called the *noise model*), and uses this model to pinpoint the ranking that is most likely to have generated the observed votes. This approach dates back to Marquis de Condorcet, who proposed a compellingly simple noise model: each voter ranks each pair of alternatives correctly with probability $p > 1/2$ and incorrectly with probability $1 - p$, and the mistakes are i.i.d. Intuitively, if a ranking is not obtained because of cycle formation, the process is restarted. Today this noise model is typically named after Mallows [139]. Probability theory was still in its infancy in the 18th century (in fact Condorcet was one of its pioneers), so the maximum likelihood estimator in Mallows' model — the *Kemeny* rule [119] — had to wait another two centuries to receive due recognition [209]. More recently, the MLE approach has received some attention in computer science [66, 83, 85, 140, 179], in part because its main prerequisite (underlying true ranking) is naturally satisfied by some of the crowdsourcing and human computation domains, where voting is in fact commonly used [140, 179].

As compelling as the MLE approach is, it has two significant drawbacks. First, it insists that the voting rule be the MLE, which is a tall order given that there is a unique MLE for the assumed noise model, and statistical accuracy may not be the only consideration in choosing a voting rule. This is reflected in existing negative results [66, 85]. To that end, we relax the requirement by asking: *How many* votes do prominent voting rules need to recover the true ranking with high probability? In crowdsourcing tasks, for example, the number of votes required directly translates to the amount of time and money one must spend to obtain accurate results. Taking one step further and adopting a more normative viewpoint, we ask: Which voting rules are *accurate in the limit*, i.e., are guaranteed to return the correct ranking given an *infinite* number of samples?

A second, important shortcoming of the MLE approach is that it relies on a specific noise model. In practice, it is extremely difficult, if not impossible, to pinpoint the exact noise model that generates votes, and classical noise models propose in the literature fall short of explaining noisy votes in the real world [140]. When the assumed noise models differs from the one generating the votes, even if there is one, the MLE rule we design would provide no guarantees. To alleviate this issue, we propose the design of *robust voting rules*, which provide guarantees with respect to a wide family of noise models, thus providing guarantees as long as the true noise model belongs to this wide family. As we cannot expect a single voting rule to be MLE for multiple noise models, we only ask the voting rule to be accurate in the limit.

7.2 Preliminaries

As this chapter focuses on a fundamentally different paradigm of social choice theory than the paradigm studied in the previous chapter, we need additional terminology. For the sake of completeness, we provide below a complete set of notations, which may have repetitions or redefinitions of some of the notations previously introduced.

We consider a set A of m alternatives. Let $\mathcal{L}(A)$ be the set of votes (which we may think of as rankings or permutations), and $\mathcal{D}(\mathcal{L}(A))$ be the set of probability distributions over $\mathcal{L}(A)$. Each vote $\sigma \in \mathcal{L}(A)$ is a bijection $\sigma : A \to \{1, 2, \ldots, m\}$. Hence, $\sigma(a)$ is the position of alternative a in σ. In particular, $\sigma(a) < \sigma(b)$ denotes that a is preferred to b under σ; we alternatively denote this by $a \succ_\sigma b$. A *vote profile* (or simply *profile*) $\pi \in \mathcal{L}(A)^n$ consists of a set of n votes for some $n \in \mathbb{N}$.

7.2.1 Voting rules

A *deterministic voting rule* is a function $r : \cup_{n \geqslant 1} \mathcal{L}(A)^n \to \mathcal{L}(A)$ which operates on a vote profile and outputs a ranking. First, note that we define the voting rule to output a ranking over alternatives rather than a single alternative; such functions are also known as *social welfare functions* (SWFs) in the literature.[1] Second, in contrast to the traditional notation, we define a voting rule to operate on any number of votes, which is required to analyze its asymptotic properties as the number of votes grows. We consider *randomized voting rules* which are denoted by $r : \cup_{n \geqslant 1} \mathcal{L}(A)^n \to \mathcal{D}(\mathcal{L}(A))$ where $D(\cdot)$ denotes the set of all distributions over an outcome space. We use $\Pr[r(\pi) = \sigma]$ to denote the probability of rule r returning ranking σ given profile π. Next, we define (families of) voting rules that play a key role in the chapter.

[1]In contrast, functions that map a vote profile to a single alternative (or a distribution over alternatives, if randomized) are known as social choice functions (SCFs).

GSRs [204]: We say that two vectors $y, z \in \mathbb{R}^k$ are *equivalent* (denoted $y \sim z$) if for every $i, j \in [k]$ we have $y_i \geqslant y_j \Leftrightarrow z_i \geqslant z_j$. We say that a function $g : \mathbb{R}^k \to \mathcal{D}(\mathcal{L}(A))$ is *compatible* if $y \sim z$ implies $g(y) = g(z)$. A *generalized scoring rule (GSR)* is given by a pair of functions (f, g), where $f : \mathcal{L}(A) \to \mathbb{R}^k$ maps every ranking to a k-dimensional vector and $g : \mathbb{R}^k \to \mathcal{D}(\mathcal{L}(A))$ is a compatible function that maps every k-dimensional vector to a distribution over rankings, and the output of the rule on a profile $\pi = (\sigma_1, \ldots, \sigma_n)$ is given by $g\left(\sum_{i=1}^n f(\sigma_i)\right)$. GSRs are characterized by two social choice axioms [205], and have interesting connections to machine learning [203]. While GSRs were originally introduced as deterministic SCFs, the definition naturally extends to (possibly randomized) SWFs.

Popular SWFs. Let us define some popular SWFs that are captured by the families of SWFs defined above. First, define the weighted pairwise majority (PM) graph of a profile as the graph where the alternatives are the vertices and there is an edge from every alternative a to every other alternative b with weight equal to the fraction of voters that prefer a to b. Its main difference from the unweighted PM graph defined above is that the latter has at most one unweighted directed edge between two alternatives a and b, indicating which alternative is preferred by the (strict) majority of the voters.

- *The Kemeny rule*: Given a profile π, the Kemeny score of a ranking is the total weight of the edges of the weighted PM graph of π in its direction. The Kemeny rule selects the ranking with the highest Kemeny score. Tie-breaking is used to choose among all the rankings with identical highest Kemeny score.

  Alternatively, given a profile $\pi = (\sigma_1, \ldots, \sigma_n) \in \mathcal{L}(A)^n$, the Kemeny rule selects a ranking $\sigma \in \mathcal{L}(A)$ that minimizes $\sum_{i=1}^n d_{KT}(\sigma, \sigma_i)$, where d_{KT} is the *Kendall tau (KT)* distance defined as

  $$d_{KT}(\sigma_1, \sigma_2) = |\{(a, b) | ((a \succ_{\sigma_1} b) \wedge (b \succ_{\sigma_2} a)) \vee ((b \succ_{\sigma_1} a) \wedge (a \succ_{\sigma_2} b))\}|.$$

  In words, the KT distance between two rankings is their number of disagreements over pairs of alternatives, and informally it is equal to the minimum number of adjacent swaps required to convert one ranking into the other. We give special attention to the *Kemeny rule with uniform tie-breaking* — the randomized version of the Kemeny rule where ties are broken uniformly, i.e., each ranking in $\arg\min_{\sigma \in \mathcal{L}(A)} \sum_{i=1}^n d_{KT}(\sigma, \sigma_i)$ is returned with equal probability.

- *(Positional) Scoring Rules* A scoring rule is given by a scoring vector $\alpha = (\alpha_1, \ldots, \alpha_m)$ where $\alpha_i \geqslant \alpha_{i+1}$ for all $i \in \{1, \ldots, m\}$ and $\alpha_1 > \alpha_m$. Under this rule for each vote σ and $i \in \{1, \ldots, m\}$, α_i points are awarded to the alternative $\sigma^{-1}(i)$, that is, α_1 points to the first alternative, α_2 points to the second alternative, and so on. The alternative with the most points overall is selected as the winner. We naturally extend this rule to output the ranking where alternatives are sorted in the descending order of their total points. Our results on positional scoring rules hold irrespective of the tie-breaking rule used. Special scoring rules include *plurality* with $\alpha = (1, 0, 0, \ldots, 0)$, *Borda count* with $\alpha = (m, m-1, \ldots, 1)$, the *veto* rule with $\alpha = (1, 1, \ldots, 1, 0)$, and the *harmonic* rule [34] with $\alpha = (1, 1/2, \ldots, 1/m)$.

- *Single transferable vote (STV)*: STV proceeds in rounds, where in each round the alternative with lowest plurality score is eliminated, until only one alternative remains. The rule ranks the alternatives in the reverse order of their elimination. At each stage, tie-breaking is used to choose among the alternatives with identical plurality score in that stage.
- *Copeland's method*: The Copeland score of an alternative a in a profile π, denoted $PW^{\pi}(a)$, is the number of outgoing edges from a in the unweighted PM graph of π, i.e., the number of other alternatives that a defeats in a pairwise election. Copeland's method ranks the alternative in non-increasing order of their Copeland scores. Tie-breaking is used for sorting alternatives with identical Copeland scores.
- *The maximin rule*: Given a profile π, the maximin score of an alternative a is the minimum of the weights of the alternative's outgoing edges in the weighted PM graph of π. The maximin rule returns the alternatives in descending order of their maximin score. Tie-breaking is used to sort alternatives with identical maximin scores.
- *The Slater rule*: Given a profile π, the Slater rule selects the ranking which minimizes the number of pairs of alternatives on which it disagrees with the unweighted PM graph of π. Note that this is, in some sense, the unweighted version of the Kemeny rule, which, as defined above, minimizes disagreement with the weighted PM graph of π. Tie-breaking is used to choose among rankings having equal disagreement with the unweighted PM graph of π.
- *The Bucklin rule*: The Bucklin score of an alternative a is the minimum k such that a is ranked among the first k positions by a majority of the voters. The Bucklin rule sorts the alternatives in a non-decreasing order according to their Bucklin scores. Tie-breaking is used to sort alternatives with identical Bucklin scores.[2]
- *The ranked pairs method*: Under the ranked pairs method, given a profile π, all ordered pairs of alternatives (a, a') are sorted in a non-increasing order of the weight of the edge from a to a' in the weighted PM graph of π. Then, starting with the first pair in the list, the method "locks in" the outcome using the result of the pairwise comparison. It proceeds with subsequent pairs and locks in every pairwise result that does not contradict (by forming a cycle) the partial ordering established so far. Finally, the method outputs the total order obtained. Tie-breaking is used initially to sort ordered pairs of alternatives with identical weight in the weighted PM graph of π.

[2]Sometimes, a deterministic scheme is used to break ties by the number of votes that rank an alternative among the first k positions, where k is the Bucklin score of the alternative. However, we cling to our assumption of an inclusive tie-breaking scheme for uniformity.

7.2.2 Noise models and distances

We assume that there exists a true hidden order $\sigma^* \in \mathcal{L}(A)$ over the alternatives. We denote the alternative at position i in σ^* by a_i, i.e., $\sigma^*(a_i) = i$. A *noise model* G is a collection of probability distributions over rankings. For every $\sigma \in \mathcal{L}(A)$, $G(\sigma)$ denotes the distribution from which noisy estimates are generated when the ground truth is σ. The probability of sampling $\sigma' \in \mathcal{L}(A)$ from this distribution is denoted by $\Pr_G[\sigma';\sigma]$.

Our noise models are parametrized by distance functions over rankings. A *distance metric* (or distance function) over $\mathcal{L}(A)$ is a function $d(\cdot,\cdot)$ that satisfies the following properties for all $\sigma, \sigma', \sigma'' \in \mathcal{L}(A)$:

- $d(\sigma,\sigma') \geqslant 0$, and $d(\sigma,\sigma') = 0$ if and only if $\sigma = \sigma'$.
- $d(\sigma,\sigma') = d(\sigma',\sigma)$.
- $d(\sigma,\sigma'') + d(\sigma'',\sigma') \geqslant d(\sigma,\sigma')$.

We assume that our distance functions are right-invariant: the distance between any two rankings does not change if the alternatives are relabeled, which is a standard assumption. A right-invariant distance function is fully specified by the distances of all rankings from any single base ranking.

We consider three popular distance functions in this chapter: the Kendall tau (KT) distance (which we have defined above), the *footrule* distance, and the *maximum displacement* distance. We investigate the KT distance (which we already defined above) in detail in Section 7.3. Definitions of the other two distance functions are given below.

The Footrule Distance The footrule distance measures the total displacements of the alternatives between two rankings. Formally, $d_{FR}(\sigma_1,\sigma_2) = \sum_{a\in A} |\sigma_1(a) - \sigma_2(a)|$.

The Maximum Displacement Distance The maximum displacement distance measures the maximum displacement of any alternative between two rankings. Formally, $d_{MD}(\sigma_1,\sigma_2) = \max_{a\in A} |\sigma_1(a) - \sigma_2(a)|$.

A noise model defines the probability of observing a ranking given an underlying true ranking, i.e., $\Pr[\sigma|\sigma^*]$ for all $\sigma, \sigma^* \in \mathcal{L}(A)$. In Section 7.3, we focus on a particular noise model, known as the *Mallows model* [139], which is widely used in machine learning and statistics. In this model, a ranking is generated given the true ranking σ^* as follows. When two alternatives a and b with $a \succ_{\sigma^*} b$ are compared, the outcome is consistent with the true ranking, i.e., $a \succ b$, with a fixed probability $1/2 < p < 1$. Every two alternatives are compared in this manner, and the process is restarted if the generated vote has a cycle (e.g., $a \succ b \succ c \succ a$). It is easy to check that the probability of drawing a ranking σ, given that the true order is σ^*, is proportional to

$$p^{\binom{m}{2} - d_{KT}(\sigma,\sigma^*)} \cdot (1-p)^{d_{KT}(\sigma,\sigma^*)},$$

which upon normalization gives

$$\Pr[\sigma|\sigma^*] = \frac{\varphi^{d_{KT}(\sigma,\sigma^*)}}{Z_\varphi^m},$$

where $\varphi = (1-p)/p < 1$ and Z_φ^m is the normalization factor which is independent of the true ranking σ^* (see, e.g., [137]). Let $p_{i,j}$ denote the probability that the alternative at position i in the true ranking (a_i) appears in position j in a random vote, so

$$p_{i,j} = \sum_{\sigma \in \mathcal{L}(A) | \sigma(a_i) = j} \Pr[\sigma | \sigma^*].$$

Let $q_{i,k} = \sum_{j=1}^{k} p_{i,j}$. Votes are sampled independently, so the probability of observing a profile $\pi = (\sigma_1, \ldots, \sigma_n) \in \mathcal{L}(A)^n$ is $\Pr[\pi|\sigma^*] = \prod_{i=1}^{n} \Pr[\sigma_i|\sigma^*]$. We note that this model is equivalent to the *Condorcet noise model*.

7.3 Sample Complexity in Mallows' Model

We first consider Mallows' model and analyze the number of samples needed by different voting rules to determine the true ranking with high probability; we use this *sample complexity* as a criterion to distinguish among voting rules or families of voting rules. For any (randomized) voting rule r, integer $k \in \mathbb{N}$ and ranking $\sigma \in \mathcal{L}(A)$, let $\mathrm{Acc}^r(k, \sigma) = \sum_{\pi \in \mathcal{L}(A)^k} \Pr[\pi|\sigma] \cdot \Pr[r(\pi) = \sigma]$ denote the accuracy of rule r with k samples and true ranking σ, that is, the probability that rule r returns σ given k samples from Mallows' model with true ranking σ. We overload the notation by letting $\mathrm{Acc}^r(k) = \min_{\sigma \in \mathcal{L}(A)} \mathrm{Acc}^r(k, \sigma)$. In words, given k samples from Mallows' model, rule r returns the underlying true ranking with probability at least $\mathrm{Acc}^r(k)$ irrespective of what the true ranking is. Finally, we denote $N^r(\varepsilon) = \min\{k \mid \mathrm{Acc}^r(k) \geqslant 1 - \varepsilon\}$, which is the number of samples required by rule r to return the true ranking with probability at least $1 - \varepsilon$. Informally, we call $N^r(\varepsilon)$ the sample complexity of rule r.

We begin by showing that for any number of alternatives m and any accuracy level ε, the Kemeny rule (with uniform tie-breaking) requires the minimum number of samples from Mallows' model to determine the true ranking with probability at least $1 - \varepsilon$. It is already known that the Kemeny rule is the maximum likelihood estimator (MLE) for the true ranking given samples from Mallows' model. Formally, given a profile $\pi = (\sigma_1, \ldots, \sigma_n)$ from Mallows' model, the MLE estimator of the true ranking is

$$\arg\max_{\sigma \in \mathcal{L}(A)} \Pr[\pi|\sigma] = \arg\max_{\sigma \in \mathcal{L}(A)} \prod_{i=1}^{n} \frac{\varphi^{d_{KT}(\sigma_i, \sigma)}}{Z_\varphi^m} = \arg\min_{\sigma \in \mathcal{L}(A)} \sum_{i=1}^{n} d_{KT}(\sigma_i, \sigma),$$

where the expression on the right hand side is a Kemeny ranking. While at first glance it may seem that this directly implies optimal sample complexity of the Kemeny rule, we give an example in Appendix B.1 of a noise model where the MLE rule does not have optimal sample complexity. However, we show that for Mallows' model, the Kemeny rule is optimal in terms of sample complexity.

Let KEM denote the Kemeny rule where ties are broken uniformly at random. That is, for any profile $\pi = (\sigma_1, \ldots, \sigma_n) \in \mathcal{L}(A)^n$, let $\text{TIE-KEM}(\pi) = \arg\min_{\sigma \in \mathcal{L}(A)} \sum_{i=1}^{n} d_{KT}(\sigma_i, \sigma)$ denote the set of Kemeny rankings. Then for every $\sigma \in \text{TIE-KEM}(\pi)$, we have $\Pr[\text{KEM}(\pi) = \sigma] = 1/|\text{TIE-KEM}(\pi)|$.

Theorem 7.1. *The Kemeny rule with uniform tie-breaking has the optimal sample complexity in Mallows' model, that is, for any number of alternatives m and any $\varepsilon > 0$, $N^{\text{KEM}}(\varepsilon) \leqslant N^r(\varepsilon)$ for every (randomized) voting rule r.*

Proof. Note that by definition of $N^r(\varepsilon)$, it is sufficient to show that the Kemeny rule has the greatest accuracy among all voting rules for any number of samples, that is, $\text{Acc}^{\text{KEM}}(k) \geqslant \text{Acc}^r(k)$ for all rules r and all $k > 0$. To show that KEM has the greatest accuracy, we need two lemmas. Define $\text{TotAcc}^r(k) = \sum_{\sigma \in \mathcal{L}(A)} \text{Acc}^r(k, \sigma)$.

Lemma 7.1. $\text{Acc}^{\text{KEM}}(k, \sigma) = \text{Acc}^{\text{KEM}}(k, \sigma'), \forall \sigma, \sigma' \in \mathcal{L}(A), \forall k \in \mathbb{N}$.

Lemma 7.2. $\text{TotAcc}^{\text{KEM}}(k) \geqslant \text{TotAcc}^r(k), \forall$ *rule* $r, \forall k \in \mathbb{N}$.

First, it is easy to derive the final result using Lemmas 7.1 and 7.2. Fix any $\varepsilon > 0$ and let $N^{\text{KEM}}(\varepsilon) = k$. Then, there exists $\widehat{\sigma} \in \mathcal{L}(A)$ such that $\text{Acc}^{\text{KEM}}(k-1, \widehat{\sigma}) < 1 - \varepsilon$, hence $\text{Acc}^{\text{KEM}}(k-1, \sigma) < 1 - \varepsilon$ for every $\sigma \in \mathcal{L}(A)$ due to Lemma 7.1. Hence, $\text{TotAcc}^{\text{KEM}}(k - 1) < m! \cdot (1 - \varepsilon)$. Now for any voting rule r, Lemma 7.2 implies $\text{TotAcc}^r(k-1) \leqslant \text{TotAcc}^{\text{KEM}}(k-1) < m! \cdot (1-\varepsilon)$, and hence by pigeonhole principle, there exists $\sigma \in \mathcal{L}(A)$ such that $\text{Acc}^r(k-1, \sigma) < 1 - \varepsilon$. Therefore, $N^r(\varepsilon) \geqslant k = N^{\text{KEM}}(\varepsilon)$, as required. Now we prove Lemmas 7.1 and 7.2.

of Lemma 7.1. Take any $k \in \mathbb{N}$ and $\sigma, \sigma' \in \mathcal{L}(A)$. Let $\omega : A \to A$ be the (unique) bijection that when applied on σ gives σ'. That is, $\omega(\sigma(i)) = \sigma'(i)$ for all $1 \leqslant i \leqslant m$. We abuse the notation and extend ω to a bijection $\omega : \mathcal{L}(A) \to \mathcal{L}(A)$ where for any $\tau \in \mathcal{L}(A)$, we have $(\omega(\tau))(i) = \omega(\tau(i))$. Essentially, we apply ω on each element of a ranking. So $\omega(\sigma) = \sigma'$. Finally, we further extend ω to operate on profiles where we apply ω to each ranking in the profile individually. Then,

$$\begin{aligned}
\text{Acc}^{\text{KEM}}(k, \sigma') &= \sum_{\pi \in \mathcal{L}(A)^k} \Pr[\pi | \sigma'] \cdot \Pr[\text{KEM}(\pi) = \sigma'] \\
&= \sum_{\omega(\pi) \in \mathcal{L}(A)^k} \Pr[\omega(\pi) | \sigma'] \cdot \Pr[\text{KEM}(\omega(\pi)) = \sigma'] \\
&= \sum_{\pi \in \mathcal{L}(A)^k} \Pr[\pi | \omega^{-1}(\sigma')] \cdot \Pr[\text{KEM}(\pi) = \omega^{-1}(\sigma')] \\
&= \sum_{\pi \in \mathcal{L}(A)^k} \Pr[\pi | \sigma] \cdot \Pr[\text{KEM}(\pi) = \sigma] = \text{Acc}^{\text{KEM}}(k, \sigma).
\end{aligned}$$

The second transition follows since ω is a bijection, the third transition follows since Mallows' model and Kemeny rule with uniform tie-breaking are anonymous with respect to the alternatives (note that uniform tie-breaking plays an important role), and the fourth transition follows since $\omega^{-1}(\sigma') = \sigma$. ■ (Lemma 7.1)

of Lemma 7.2. For any rule r and any $k \in \mathbb{N}$,

$$\text{TotAcc}^r(k) = \sum_{\sigma \in \mathcal{L}(A)} \sum_{\pi \in \mathcal{L}(A)^k} \Pr[\pi | \sigma] \cdot \Pr[r(\pi) = \sigma]$$

$$
\begin{aligned}
&= \sum_{\pi \in \mathcal{L}(A)^k} \sum_{\sigma \in \mathcal{L}(A)} \Pr[\pi|\sigma] \cdot \Pr[r(\pi) = \sigma] \\
&\leqslant \sum_{\pi \in \mathcal{L}(A)^k} \sum_{\sigma \in \mathcal{L}(A)} \Pr[r(\pi) = \sigma] \cdot \left(\max_{\sigma' \in \mathcal{L}(A)} \Pr[\pi|\sigma'] \right) \\
&= \sum_{\pi \in \mathcal{L}(A)^k} \max_{\sigma' \in \mathcal{L}(A)} \Pr[\pi|\sigma'] \\
&= \sum_{\pi \in \mathcal{L}(A)^k} \max_{\sigma' \in \mathcal{L}(A)} \Pr[\pi|\sigma'] \cdot \sum_{\sigma \in \textsc{Tie-Kem}(\pi)} \frac{1}{|\textsc{Tie-Kem}(\pi)|} \\
&= \sum_{\pi \in \mathcal{L}(A)^k} \sum_{\sigma \in \textsc{Tie-Kem}(\pi)} \Pr[\pi|\sigma] \cdot \Pr[\mathrm{KEM}(\pi) = \sigma] \\
&= \mathrm{TotAcc}^{\mathrm{KEM}}(k),
\end{aligned}
$$

where the sixth transition holds because the Kemeny rule is an MLE for Mallows' model, hence $\max_{\sigma' \in \mathcal{L}(A)} \Pr[\pi|\sigma'] = \Pr[\pi|\sigma]$ for every $\sigma \in \textsc{Tie-Kem}(\pi)$. Also, under uniform tie-breaking $\Pr[\mathrm{KEM}(\pi) = \sigma] = 1/|\textsc{Tie-Kem}(\pi)|$ for every $\sigma \in \textsc{Tie-Kem}(\pi)$. ■ (Lemma 7.2)

Thus, we have established that to output the true underlying ranking with any given probability, the Kemeny rule with uniform tie-breaking requires the minimum number of samples from Mallows' model among all voting rules. ■ (Theorem 7.1)

Now that we know that the Kemeny rule has the optimal sample complexity, a natural question is to determine how many samples it really requires. Instead of analyzing the sample complexity of the Kemeny rule particularly, we consider a family of voting rules (which includes the Kemeny rule itself) such that each rule in this family has the same asymptotic sample complexity as that of the Kemeny rule.

7.3.1 The family of PM-c rules

Our family of voting rules crucially relies on the standard concept of *pairwise-majority graph* (PM graph). Given a profile $\pi \in \mathcal{L}(A)^n$, the PM graph of π is the directed graph $G = (V, E)$, where the alternatives are the vertices ($V = A$) and there is an edge from alternative a to alternative a' if a is preferred to a' in a (strong) majority of the rankings of π. Formally, $(a, a') \in E$ if $|\{\sigma \in \pi | a \succ_\sigma a'\}| > |\{\sigma \in \pi | a' \succ_\sigma a\}|$. Note that there may be pairs of alternatives such that there is no edge in the PM graph in either direction (if they are tied), but there can never be an edge in both directions. A PM graph can also have directed cycles. When a PM graph is complete (i.e., there is an edge between every pair of alternatives) and acyclic, there exists a unique $\sigma \in \mathcal{L}(A)$ such that there is an edge from a to a' if and only if $a \succ_\sigma a'$. In this case, we say that the PM graph reduces to σ.

Definition 7.1 (Pairwise-Majority Consistent Rules)**.** A deterministic voting rule r is called pairwise-majority consistent (PM-c) if $r(\pi) = \sigma$ whenever the PM graph of π reduces to σ.[3] For randomized voting rules, we require that $\Pr[r(\pi) = \sigma] = 1$.

To the best of our knowledge this family of rules is novel. Note though that an acyclic and complete PM graph is similar to — and in some sense an extension of — having a Condorcet winner. A Condorcet winner is an alternative that beats every other alternative in a pairwise election. It is easy to check that if such an alternative exists, then it is unique and it is a source in the PM graph with $m-1$ outgoing edges and no incoming edges. Thus, profiles where the PM graph reduces to a ranking necessarily have a Condorcet winner. In addition, they have a second alternative with $m-2$ outgoing edges and only 1 incoming edge, a third alternative with $m-3$ outgoing edges and 2 incoming edges, and so on.

Theorem 7.2. *The Kemeny rule, the Slater rule, the ranked pairs method, Copeland's method, and Schulze's method are PM-c.*

The definitions of these rules and the proof of the theorem appear in Appendix B.2. Note that all the rules in Theorem 7.2 are Condorcet consistent[4] when they output a single alternative. If we take any Condorcet consistent method, apply it on a profile, remove the winner from every vote in the profile, apply the method again on the reduced profile, and keep repeating this process, then the extended rule that outputs the alternatives in the order of removal is always a PM-c rule. However, Copeland's method in Theorem 7.2 is extended by sorting the alternatives by their Copeland scores.

We now proceed to prove that any PM-c rule requires at most a logarithmic number of samples in m (the number of alternatives) and $1/\varepsilon$ to determine the true ranking with probability at least $1-\varepsilon$. First, we introduce a property of distance functions that will be used throughout the chapter. For any $\sigma \in \mathcal{L}(A)$ and $a, b \in A$, define $\sigma_{a\leftrightarrow b}$ to be the ranking obtained by swapping a and b in σ. That is, $\sigma_{a\leftrightarrow b}(c) = \sigma(c)$ for any $c \in A \setminus \{a, b\}$, $\sigma_{a\leftrightarrow b}(a) = \sigma(b)$ and $\sigma_{a\leftrightarrow b}(b) = \sigma(a)$.

Definition 7.2 (Swap-Increasing Distance Functions)**.** An integer-valued distance function d is called swap-increasing if for any $\sigma^*, \sigma \in \mathcal{L}(A)$ and alternatives $a, b \in A$ such that $a \succ_{\sigma^*} b$ and $a \succ_\sigma b$, we have $d(\sigma_{a\leftrightarrow b}, \sigma^*) \geqslant d(\sigma, \sigma^*) + 1$, and if $\sigma^*(b) = \sigma^*(a) + 1$ (a and b are adjacent in σ^*) then $d(\sigma_{a\leftrightarrow b}, \sigma^*) = d(\sigma, \sigma^*) + 1$.

The following lemma is a folklore result; we reconstruct its proof for the sake of completeness.

Lemma 7.3. *The Kendall tau (KT) distance is swap-increasing.*

Proof. Let $\sigma^*, \sigma \in \mathcal{L}(A)$ and $a, b \in A$ with $a \succ_{\sigma^*} b$ and $a \succ_\sigma b$. Let $\sigma(a) = i$ and $\sigma(b) = j$, so $i < j$. Define $Y = \{y \in A | i < \sigma(y) < j\}$. Since $\sigma^*(a) < \sigma^*(b)$, the following properties hold:

[3]In this case, ranking σ is sometimes called the "Condorcet order" of profile π.

[4]Given a profile, an alternative is called Condorcet winner if it defeats every other alternative in a pairwise election. A Condorcet winner, if one exists, is unique. A voting rule that returns a single alternative is called Condorcet consistent if it returns the Condorcet winner on every profile that admits one.

1. For every $y \in Y$, $\sigma^*(y) < \sigma^*(a)$ implies that $\sigma^*(y) < \sigma^*(b)$. Hence,
$$\sum_{y \in Y} \mathbb{1}[\sigma^*(y) < \sigma^*(a)] \leqslant \sum_{y \in Y} \mathbb{1}[\sigma^*(y) < \sigma^*(b)].$$
2. For every $y \in Y$, $\sigma^*(b) < \sigma^*(y)$ implies that $\sigma^*(a) < \sigma^*(y)$. Hence,
$$\sum_{y \in Y} \mathbb{1}[\sigma^*(b) < \sigma^*(y)] \leqslant \sum_{y \in Y} \mathbb{1}[\sigma^*(a) < \sigma^*(y)].$$
3. $\mathbb{1}[\sigma^*(a) < \sigma^*(b)] = 1$.
4. $\mathbb{1}[\sigma^*(b) < \sigma^*(a)] = 0$.

Now, we can express $d_{KT}(\sigma_{a \leftrightarrow b}, \sigma^*) - d_{KT}(\sigma, \sigma^*)$ as

$$\begin{aligned} &d_{KT}(\sigma_{a \leftrightarrow b}, \sigma^*) - d_{KT}(\sigma, \sigma^*) \\ &\quad = \sum_{y \in Y} \mathbb{1}[\sigma^*(y) < \sigma^*(b)] + \sum_{y \in Y} \mathbb{1}[\sigma^*(a) < \sigma^*(y)] + \mathbb{1}[\sigma^*(a) < \sigma^*(b)] \\ &\qquad - \sum_{y \in Y} \mathbb{1}[\sigma^*(y) < \sigma^*(a)] - \sum_{y \in Y} \mathbb{1}[\sigma^*(b) < \sigma^*(y)] - \mathbb{1}[\sigma^*(b) < \sigma^*(a)] \\ &\quad \geqslant 1, \end{aligned}$$

as desired. When a and b are adjacent in σ^* (i.e., $\sigma^*(b) = \sigma^*(a) + 1$), we show that equality holds. In this case, observe that the implications in properties (1) and (2) are actually equivalences and the inequalities can be replaced by equalities. Then, the sums in the above derivation cancel out, and it can be seen that the distance increases by exactly 1. ■ (Lemma 7.3)

We are now ready to analyze the sample complexity of PM-c rules.

Theorem 7.3. *For any given $\varepsilon > 0$, any PM-c rule determines the true ranking with probability at least $1 - \varepsilon$ given $O(\ln(m/\varepsilon))$ samples from Mallows' model.*

Proof. Let σ^* denote the true underlying ranking. We show that the PM graph of a profile of $O(\ln(m/\varepsilon))$ votes from Mallows' model reduces to σ^* with probability at least $1 - \varepsilon$. It follows that any PM-c rule would output σ^* with probability at least $1 - \varepsilon$.

Let $\pi \in \mathcal{L}(A)^n$ denote a profile of n samples from Mallows' model. For any $a, b \in A$, let n_{ab} denote the number of rankings $\sigma \in \pi$ such that $a \succ_\sigma b$. Hence, $n_{ab} + n_{ba} = n$ for every $a, b \in A$. The PM graph of π reduces to σ^* if and only if for every $a, b \in A$ such that $a \succ_{\sigma^*} b$, we have $n_{ab} - n_{ba} \geqslant 1$. Hence, we want an upper bound on n such that

$$\Pr\left[\forall a, b \in A, a \succ_{\sigma^*} b \Rightarrow n_{ab} - n_{ba} \geqslant 1\right] \geqslant 1 - \varepsilon.$$

For any $a, b \in A$ with $a \succ_{\sigma^*} b$, define $\delta_{ab} = \mathbb{E}[(n_{ab} - n_{ba})/n]$. Let $p_{a \succ b}$ denote the probability that $a \succ_\sigma b$ in a random sample σ. Then, by linearity of expectation, we have $\delta_{ab} = p_{a \succ b} - p_{b \succ a}$. Thus,

$$\Pr\left[n_{ab} - n_{ba} \leqslant 0\right] = \Pr\left[\frac{n_{ab} - n_{ba}}{n} \leqslant 0\right] \leqslant \Pr\left[\left|\frac{n_{ab} - n_{ba}}{n} - \mathbb{E}\left[\frac{n_{ab} - n_{ba}}{n}\right]\right| \geqslant \delta_{ab}\right]$$

$$\leqslant 2 \cdot e^{-2\cdot\delta_{ab}^2\cdot n} \leqslant 2 \cdot e^{-2\cdot\delta_{\min}^2\cdot n},$$

where the third transition is due to Hoeffding's inequality and in the last transition $\delta_{\min} = \min_{a,b\in A: a\succ_{\sigma^*} b} \delta_{ab}$. Applying the union bound, we get

$$\Pr\left[\exists a,b \in A,\ \{(a \succ_{\sigma^*} b) \wedge (n_{ab} - n_{ba} \leqslant 0)\}\right] \leqslant \binom{m}{2} \cdot 2 \cdot e^{-2\cdot\delta_{\min}^2\cdot n} \leqslant m^2 \cdot e^{-2\cdot\delta_{\min}^2\cdot n}.$$

It is easy to check that the right-most quantity above is at most ε when $n \geqslant \frac{1}{2\cdot\delta_{\min}^2} \cdot \ln\left(\frac{m^2}{\varepsilon}\right)$. To complete the proof we need to show that $\delta_{\min} = \Omega(1)$, that is, it is lower bounded by a constant independent of m. This is quite intuitive since the process of generating a sample from Mallows' model maintains the order between every pair of alternatives with a constant probability $p > 1/2$. However, the fact that we restart the process if a cycle is formed makes the probabilities as well as this analysis a bit more involved. For any $a, b \in A$ such that $a \succ_{\sigma^*} b$, we have

$$\begin{aligned}
\delta_{ab} &= p_{a\succ b} - p_{b\succ a} \\
&= \sum_{\sigma\in L(A)|a\succ_\sigma b} \Pr[\sigma|\sigma^*] - \sum_{\sigma\in L(A)|b\succ_\sigma a} \Pr[\sigma|\sigma^*] \\
&= \sum_{\sigma\in L(A)|a\succ_\sigma b} \left(\Pr[\sigma|\sigma^*] - \Pr[\sigma_{a\leftrightarrow b}|\sigma^*]\right) \\
&= \sum_{\sigma\in L(A)|a\succ_\sigma b} \frac{\varphi^{d_{KT}(\sigma,\sigma^*)} - \varphi^{d_{KT}(\sigma_{a\leftrightarrow b},\sigma^*)}}{Z_\varphi^m} \\
&\geqslant \sum_{\sigma\in L(A)|a\succ_\sigma b} \frac{\varphi^{d_{KT}(\sigma,\sigma^*)} \cdot (1-\varphi)}{Z_\varphi^m} \\
&= (1-\varphi) \cdot p_{a\succ b} = (1-\varphi) \cdot \left(\frac{1+\delta_{ab}}{2}\right),
\end{aligned} \tag{7.1}$$

where the third transition holds because $\sigma \leftrightarrow \sigma_{a\leftrightarrow b}$ is a bijection between all rankings where $a \succ b$ and all rankings where $b \succ a$, the fifth transition follows by using $\varphi < 1$ and Lemma 7.3, and the last transition follows from the equalities $p_{a\succ b} - p_{b\succ a} = \delta_{ab}$ and $p_{a\succ b} + p_{b\succ a} = 1$. Solving Equation (7.1), we get $\delta_{ab} \geqslant (1-\varphi)/(1+\varphi)$ for all $a, b \in A$ with $a \succ_{\sigma^*} b$. Hence, $\delta_{\min} \geqslant (1-\varphi)/(1+\varphi)$, as required.[5] ■ (Theorem 7.3)

We have seen that PM-c rules have logarithmic sample complexity; it turns out that no rule can do better, i.e., we prove a matching lower bound that holds for any randomized voting rule.

Theorem 7.4. *For any $\varepsilon \in (0, 1/2]$, any (randomized) voting rule requires $\Omega(\ln(m/\varepsilon))$ samples from Mallows' model to determine the true ranking with probability at least $1-\varepsilon$.*

[5]Recall that $\varphi = (1-p)/p$, where $p > 1/2$ is the pairwise comparison probability under Mallows' model. Hence, we have proved $\delta_{\min} \geqslant 2 \cdot p - 1$.

Proof. Consider any voting rule r. Assume $\mathrm{Acc}^r(n) \geqslant 1-\varepsilon$ for some $n \in \mathbb{N}$. We want to show that $n = \Omega(\ln(m/\varepsilon))$. For any $\sigma \in \mathcal{L}(A)$, we have $\mathrm{Acc}^r(n,\sigma) \geqslant 1-\varepsilon$. Consider an arbitrary $\sigma \in \mathcal{L}(A)$, and let $\mathcal{N}(\sigma) = \{\sigma' \in \mathcal{L}(A) | d_{KT}(\sigma',\sigma) = 1\}$ denote the set of all rankings at distance 1 from σ. Then, for any ranking $\sigma' \in \mathcal{N}(\sigma)$ and any profile $\pi = (\sigma_1,\ldots,\sigma_n) \in \mathcal{L}(A)^n$, we have

$$\Pr[\pi|\sigma] = \prod_{i=1}^{n} \frac{\varphi^{d_{KT}(\sigma_i,\sigma)}}{Z_\varphi^m} \geqslant \prod_{i=1}^{n} \frac{\varphi^{d_{KT}(\sigma_i,\sigma')+1}}{Z_\varphi^m} = \varphi^n \cdot \Pr[\pi|\sigma'], \tag{7.2}$$

where the second transition holds since for any $\tau \in \mathcal{L}(A)$ the triangle inequality implies

$$d_{KT}(\tau,\sigma) \leqslant d_{KT}(\tau,\sigma') + d_{KT}(\sigma,\sigma') = d_{KT}(\tau,\sigma') + 1.$$

Now,

$$\begin{aligned}
\mathrm{Acc}^r(n,\sigma) &= \sum_{\pi\in\mathcal{L}(A)^n} \Pr[\pi|\sigma] \cdot \Pr[r(\pi) = \sigma] \\
&= \sum_{\pi\in\mathcal{L}(A)^n} \Pr[\pi|\sigma] \cdot (1 - \Pr[r(\pi) \neq \sigma]) \\
&= 1 - \sum_{\pi\in\mathcal{L}(A)^n} \Pr[\pi|\sigma] \cdot \Pr[r(\pi) \neq \sigma] \\
&\leqslant 1 - \sum_{\pi\in\mathcal{L}(A)^n} \Pr[\pi|\sigma] \cdot \left(\sum_{\sigma'\in\mathcal{N}(\sigma)} \Pr[r(\pi) = \sigma'] \right) \\
&\leqslant 1 - \sum_{\sigma'\in\mathcal{N}(\sigma)} \sum_{\pi\in\mathcal{L}(A)^n} \varphi^n \cdot \Pr[\pi|\sigma'] \cdot \Pr[r(\pi) = \sigma'] \\
&= 1 - \varphi^n \cdot \sum_{\sigma'\in\mathcal{N}(\sigma)} \mathrm{Acc}^r(n,\sigma') \\
&\leqslant 1 - \varphi^n \cdot (m-1) \cdot (1-\varepsilon),
\end{aligned}$$

where the fifth transition follows by changing the order of summation and Equation (7.2), and the last transition holds since $\mathrm{Acc}^r(n) \geqslant 1-\varepsilon$ and $|\mathcal{N}(\sigma)| = m-1$. Thus, to achieve an accuracy of at least $1-\varepsilon$, we need $\varphi^n \cdot (m-1) \cdot (1-\varepsilon) \leqslant \varepsilon$, and the theorem follows by solving for n. ■ (Theorem 7.4)

7.3.2 Scoring rules may require exponentially many samples

While Theorems 7.3 and 7.4 show that every PM-c rule requires an asymptotically optimal (and in particular, logarithmic) number of samples to determine the true ranking with high probability, some classical voting rules such as plurality fall short. In particular, we establish that plurality requires at least exponentially many samples to determine the true ranking with high probability. Since plurality relies on the number of appearances of various alternatives in the first position, our analysis crucially relies on the probability of different alternatives coming first in a random vote, i.e., $p_{i,1}$.

Lemma 7.4. $p_{i,1} = \varphi^{i-1} / \left(\sum_{j=1}^{m} \varphi^{j-1}\right)$ *for all* $i \in \{1, \ldots, m\}$.

Proof. Recall that a_i denotes the alternative at position i in the true ranking σ^*. First we prove that for any $i \in \{1, \ldots, m-1\}$, we have $p_{i+1,1} = \varphi \cdot p_{i,1}$. To see this,

$$\begin{aligned} p_{i,1} - p_{i+1,1} &= \frac{\sum_{\sigma \in \mathcal{L}(A) | \sigma(a_i)=1} \varphi^{d_{KT}(\sigma,\sigma^*)} - \sum_{\sigma \in \mathcal{L}(A) | \sigma(a_{i+1})=1} \varphi^{d_{KT}(\sigma,\sigma^*)}}{Z_\varphi^m} \\ &= \frac{\sum_{\sigma \in \mathcal{L}(A) | \sigma(a_i)=1} \left(\varphi^{d_{KT}(\sigma,\sigma^*)} - \varphi^{d_{KT}(\sigma_{a_i \leftrightarrow a_{i+1}},\sigma^*)}\right)}{Z_\varphi^m} \\ &= \sum_{\sigma \in \mathcal{L}(A) | \sigma(a_i)=1} \frac{\varphi^{d_{KT}(\sigma,\sigma^*)} \cdot (1-\varphi)}{Z_\varphi^m} = (1-\varphi) \cdot p_{i,1}, \end{aligned}$$

where the second transition follows since $\sigma \leftrightarrow \sigma_{a_i \leftrightarrow a_{i+1}}$ is a bijection between the set of all rankings where a_i is first and the set of all rankings where a_{i+1} is first, and the third transition follows due to Lemma 7.3. Hence, $p_{i,1} - p_{i+1,1} = (1-\varphi) \cdot p_{i,1}$, which implies that $p_{i+1,1} = \varphi \cdot p_{i,1}$. Applying this repeatedly, we have that $p_{i,1} = p_{1,1} \cdot \varphi^{i-1}$, for every $i \in \{1, \ldots, m\}$. Summing over $1 \leqslant i \leqslant m$ and observing that $\sum_{i=1}^{m} p_{i,1} = 1$, we get the desired result. ■(Lemma 7.4)■

Lemma 7.4 directly implies that the probability of sampling votes in which a_{m-1} or a_m (the two alternatives that are ranked at the bottom of σ^*) are at the top is exponentially small, hence plurality requires an exponential number of samples to "see" these alternatives and distinguish between them. What makes the proof more difficult is that in theory the tie-breaking scheme may help plurality return the true ranking; indeed it is known that the choice of tie breaking scheme can significantly affect a rule's performance [158]. However, we show that here this is not the case, i.e., our lower bound works for any natural (randomized) tie-breaking scheme.

Theorem 7.5. *For any* $\varepsilon \in (0, 1/4]$, *plurality (with any possibly randomized tie-breaking scheme that depends on the top-ranked alternatives of the input votes) requires* $\Omega((1/\varphi)^m)$ *samples from Mallows' model to output the true ranking with probability at least* $1 - \varepsilon$.

Proof. We first note that instead of operating on a profile $\pi \in \mathcal{L}(A)^n$, plurality (and its tie-breaking scheme) operates on the vector of its plurality votes $v \in A^n$ (we call it a top-vote) which consists of the top-ranked alternatives of the different votes of π. The probability of observing a top-vote v given a true ranking σ^* is the sum of the probabilities of observing profiles whose top-vote is v; we denote this by $\Pr[v|\sigma^*]$. The accuracy of the plurality rule (denoted PL) with n samples on a true ranking σ can now equivalently be written as

$$\mathrm{Acc}^{\mathrm{PL}}(n, \sigma) = \sum_{v \in A^n} \Pr[v|\sigma] \cdot \Pr[\mathrm{PL}(v) = \sigma]. \tag{7.3}$$

Fix $\varepsilon \in (0, 1/4]$ and suppose we have $\mathrm{Acc}^{\mathrm{PL}}(n) \geqslant 1 - \varepsilon$, i.e., $\mathrm{Acc}^{\mathrm{PL}}(n, \sigma) \geqslant 1 - \varepsilon$ for all $\sigma \in \mathcal{L}(A)$. We want to show that $n = \Omega((1/\varphi)^m)$. Let the set of alternatives be

$A = \{a_1, \ldots, a_m\}$. Consider two distinct rankings: $\sigma_1 = (a_1 \succ \ldots \succ a_{m-2} \succ a_{m-1} \succ a_m)$ and $\sigma_2 = (a_1 \succ \ldots \succ a_{m-2} \succ a_m \succ a_{m-1})$ (where the last two alternatives are swapped compared to σ_1). Let $\widehat{A} = A \setminus \{a_{m-1}, a_m\}$. We can decompose Equation (7.3) into two parts: (i) a summation over $v \in \widehat{A}^n$ (when plurality does not "see" alternatives a_{m-1} and a_m); denote this by $f(\sigma)$, and (ii) a summation over $v \in A^n \setminus \widehat{A}^n$ (when plurality "sees" at least one of them); denote this by $g(\sigma)$.

For any $v \in \widehat{A}^n$, we have $\Pr[v|\sigma_1] = \Pr[v|\sigma_2]$. To see this, observe that in any profile π with top-vote v we can swap alternatives a_{m-1} and a_m in all the votes to obtain (the unique) profile π' which importantly also has top-vote v and $\Pr[\pi|\sigma_1] = \Pr[\pi'|\sigma_2]$. Summing over all profiles with top-vote v, this yields $\Pr[v|\sigma_1] = \Pr[v|\sigma_2]$. Therefore, we have

$$f(\sigma_1) + f(\sigma_2) = \sum_{v \in \widehat{A}^n} \Pr[v|\sigma_1] \cdot (\Pr[\mathrm{PL}(v) = \sigma_1] + \Pr[\mathrm{PL}(v) = \sigma_2]) \leqslant \sum_{v \in \widehat{A}^n} \Pr[v|\sigma_1] \leqslant 1.$$

Further,

$$g(\sigma_1) = \sum_{v \in A^n \setminus \widehat{A}^n} \Pr[v|\sigma_1] \cdot \Pr[\mathrm{PL}(v) = \sigma_1] \leqslant \sum_{v \in A^n \setminus \widehat{A}^n} \Pr[v|\sigma_1],$$

where the right hand side is the probability that at least one of the two alternatives a_{m-1} and a_m comes first in at least one vote. Let $t_{i,j}$ denote the number of votes in which alternative a_i appears in position j. Then we have

$$g(\sigma_1) \leqslant \Pr[(t_{m-1,1} > 0) \vee (t_{m,1} > 0)] \leqslant \Pr[t_{m-1,1} > 0] + \Pr[t_{m,1} > 0],$$

where the last transition is due to the union bound.

The probability that alternative a_{m-1} appears first in a vote is $p_{m-1,1}$. Therefore, the probability that it appears first in at least one vote is at most $n \cdot p_{m-1,1}$ by the union bound. Similarly, $\Pr[t_{m,1} > 0] \leqslant n \cdot p_{m,1}$. Therefore, $g(\sigma_1) \leqslant n \cdot (p_{m-1,1} + p_{m,1})$. In the same way, we can obtain $g(\sigma_2) \leqslant n \cdot (p_{m-1,1} + p_{m,1})$. Finally, using the bounds obtained on f and g, we have

$$\mathrm{Acc}^{\mathrm{PL}}(n, \sigma_1) + \mathrm{Acc}^{\mathrm{PL}}(n, \sigma_2) = (f(\sigma_1) + f(\sigma_2)) + g(\sigma_1) + g(\sigma_2) \leqslant 1 + 2 \cdot n \cdot (p_{m-1,1} + p_{m,1}).$$

We assumed that $\mathrm{Acc}^{\mathrm{PL}}(n, \sigma) \geqslant 1 - \varepsilon$ for every $\sigma \in \mathcal{L}(A)$. Therefore, we need $1 + 2 \cdot n \cdot (p_{m-1,1} + p_{m,1}) \geqslant 2 \cdot (1 - \varepsilon)$, i.e.,

$$n \geqslant \frac{1 - 2 \cdot \varepsilon}{2 \cdot (p_{m-1,1} + p_{m,1})} \geqslant \frac{1}{8 \cdot p_{m-1,1}} = \frac{\sum_{j=0}^{m-1} \varphi^j}{8 \cdot \varphi^{m-2}} \geqslant \frac{1}{8 \cdot \varphi^{m-2}},$$

where the second transition follows since $\varepsilon \in (0, 1/4]$ and $p_{m,1} < p_{m-1,1}$, and the third transition follows by Lemma 7.4. Thus, plurality requires $\Omega((1/\varphi)^m)$ samples to output the true ranking with high probability. ■ (Theorem 7.5)

Since the exponential lower bound for plurality in Theorem 7.5 is missing a dependence on ε, it is in general incomparable to the logarithmic upper bound of PM-c rules in Theorem 7.3. However, the current bounds do show that plurality requires doubly exponentially more samples (asymptotically in m) compared to PM-c rules for any fixed $\varepsilon \in (0, 1/4]$. Plurality has terrible performance because it ranks alternatives by just observing their number of appearances in the first positions of the input votes. Similarly, consider the veto rule that essentially ranks alternatives in the ascending order of their number of appearances at the bottom of input votes. By symmetry we have $p_{m,m} = p_{1,1}$ and $p_{m-1,m} = p_{2,1}$, both of which are lower bounded by constants due to Lemma 7.4. Hence, veto requires only constantly many samples to distinguish between a_{m-1} and a_m. Nevertheless, it is difficult for both plurality and veto to distinguish between alternatives $a_{m/2}$ and $a_{m/2+1}$ that are far from both ends. Certain scoring rules, such as the Borda count or the harmonic scoring rule, take into consideration the number of appearances of an alternative at *all* positions. We show that a positional scoring rule that gives different weights to all positions and does not give some position exponentially higher weight than any other position would require only polynomially many samples.

Theorem 7.6. *Consider a positional scoring rule r with scoring vector $(\alpha_1, \ldots, \alpha_m)$. For $i \in \{1, \ldots, m-1\}$, define $\beta_i = \alpha_i - \alpha_{i+1}$. Let $\beta_{\max} = \max_{i<m} \beta_i$ and $\beta_{\min} = \min_{i<m} \beta_i$. Assume $\beta_{\min} > 0$ and let $\beta^* = \beta_{\max}/\beta_{\min}$. Then for any $\varepsilon > 0$, rule r requires $O((\beta^*)^2 \cdot m^2 \cdot \ln(m/\varepsilon))$ samples from Mallows' model to output the true ranking with probability at least $1 - \varepsilon$.*

Proof. Recall that a_i denotes the ith alternative in the true ranking σ^*. Consider a profile π consisting of n samples from Mallows' model. Let $t_{i,j}$ denote the number of times a_i appears in position j, and let $s_{i,k} = \sum_{j=1}^{k} t_{i,j}$. First, we note that for any $i \in \{1, \ldots, m\}$,

$$\begin{aligned}
\sum_{k=1}^{m-1} \beta_k \cdot s_{i,k} &= \sum_{k=1}^{m-1} \beta_k \cdot \left(\sum_{j=1}^{k} t_{i,j} \right) = \sum_{j=1}^{m-1} \left(\sum_{k=j}^{m-1} \beta_k \right) \cdot t_{i,j} = \sum_{j=1}^{m-1} (\alpha_j - \alpha_m) \cdot t_{i,j} \\
&= \sum_{j=1}^{m-1} \alpha_j \cdot t_{i,j} - \alpha_m \cdot (n - t_{i,m}) = \sum_{j=1}^{m} \alpha_j \cdot t_{i,j} - n \cdot \alpha_m,
\end{aligned}$$

where the second transition follows by switching the order of summation and the fourth transition holds because $\sum_{j=1}^{m} t_{i,j} = n$ as the total number of appearances of a_i equals the number of votes. Since $n \cdot \alpha_m$ is independent of the alternative, we can equivalently consider $\sum_{k=1}^{m-1} \beta_k \cdot s_{i,k}$ as the score of alternative a_i. Hence, for rule r to output σ^* with high probability we require $\Pr[\forall i \in \{1, \ldots, m\}, \sum_{k=1}^{m-1} \beta_k \cdot (s_{i,k} - s_{i+1,k}) > 0] \geqslant 1 - \varepsilon$. If we had $\Pr[\sum_{k=1}^{m-1} \beta_k \cdot (s_{i,k} - s_{i+1,k}) \leqslant 0] \leqslant \varepsilon/m$ for every $i \in \{1, \ldots, m\}$, then we would obtain (using the union bound) that r outputs σ^* with probability at least $1 - \varepsilon$. Observe that

$$\Pr\left[\sum_{k=1}^{m-1} \beta_k (s_{i,k} - s_{i+1,k}) \leqslant 0 \right] \leqslant e^{-\frac{2n\left(\sum_{k=1}^{m-1} \beta_j (q_{i,k} - q_{i+1,k})\right)^2}{4m^2 \beta_{\max}^2}} \leqslant e^{-\frac{n\beta_{\min}^2 \left(\sum_{k=1}^{m-1} (q_{i,k} - q_{i+1,k})\right)^2}{2m^2 \beta_{\max}^2}},$$

where $q_{i,k} = \sum_{j=1}^{k} p_{i,j}$ and the first transition follows from Hoeffding's inequality. Therefore, for this probability to be at most ε/m, it is sufficient to have

$$n \geqslant \frac{2 \cdot m^2 \cdot \beta_{\max}^2}{\beta_{\min}^2 \cdot \left(\sum_{k=1}^{m-1} (q_{i,k} - q_{i+1,k})\right)^2} \ln(m/\varepsilon).$$

Now we only need to prove that the term $\sum_{k=1}^{m-1} (q_{i,k} - q_{i+1,k})$ in the denominator is lower-bounded by a constant independent of m and ε. Note that $\sum_{k=1}^{m-1} q_{i,k} = \sum_{k=1}^{m} q_{i,k} - 1 = \sum_{k=1}^{m} \sum_{j=1}^{k} p_{i,j} - 1 = \sum_{j=1}^{m} (m - j + 1) \cdot p_{i,j} - 1 = \mathbb{E}[Borda(a_i)] - 1$, where $\mathbb{E}[Borda(a_i)]$ denotes the expected Borda score of alternative a_i under one random sample from Mallows' model. Similarly, $\sum_{k=1}^{m-1} q_{i+1,k} = \mathbb{E}[Borda(a_{i+1})] - 1$. Therefore,

$$\begin{aligned}
&\sum_{k=1}^{m-1} (q_{i,k} - q_{i+1,k}) \\
&= \mathbb{E}[Borda(a_i)] - \mathbb{E}[Borda(a_{i+1})] = \mathbb{E}[Borda(a_i) - Borda(a_{i+1})] \\
&= \sum_{\sigma \in \mathcal{L}(A)} \Pr[\sigma|\sigma^*] \cdot ((m + 1 - \sigma(a_i)) - (m + 1 - \sigma(a_{i+1}))) \\
&= \sum_{\substack{\sigma \in \mathcal{L}(A) \\ s.t.\ a_i \succ_\sigma a_{i+1}}} \Pr[\sigma|\sigma^*] \cdot (\sigma(a_{i+1}) - \sigma(a_i)) + \sum_{\substack{\sigma \in \mathcal{L}(A) \\ s.t.\ a_{i+1} \succ_\sigma a_i}} \Pr[\sigma|\sigma^*] \cdot (\sigma(a_{i+1}) - \sigma(a_i)) \\
&= \sum_{\substack{\sigma \in \mathcal{L}(A) \\ s.t.\ a_i \succ_\sigma a_{i+1}}} (\Pr[\sigma|\sigma^*] - \Pr[\sigma_{a_i \leftrightarrow a_{i+1}}|\sigma^*]) \cdot (\sigma(a_{i+1}) - \sigma(a_i)) \\
&\geqslant \sum_{\substack{\sigma \in \mathcal{L}(A) \\ s.t.\ a_i \succ_\sigma a_{i+1}}} (1 - \varphi) \cdot \Pr[\sigma|\sigma^*] \cdot 1 = (1 - \varphi) \cdot \Pr[a_i \succ a_{i+1}|\sigma^*] \geqslant 0.5 \cdot (1 - \varphi).
\end{aligned}$$

The second transition follows from the linearity of expectation. The fifth transition is true because under the bijective mapping $\sigma \leftrightarrow \sigma_{a_i \leftrightarrow a_{i+1}}$, we have $\sigma_{a_i \leftrightarrow a_{i+1}}(a_i) = \sigma(a_{i+1})$ and $\sigma_{a_i \leftrightarrow a_{i+1}}(a_{i+1}) = \sigma(a_i)$. For the sixth transition, note that in any σ where $a_i \succ_\sigma a_{i+1}$, $\sigma(a_{i+1}) \geqslant \sigma(a_i) + 1$. Also, Lemma 7.3 implies that $d_{KT}(\sigma_{a_i \leftrightarrow a_{i+1}}, \sigma^*) = d_{KT}(\sigma, \sigma^*) + 1$, so

$$\begin{aligned}
\Pr[\sigma|\sigma^*] - \Pr[\sigma_{a_i \leftrightarrow a_{i+1}}|\sigma^*] &= (\varphi^{d_{KT}(\sigma,\sigma^*)} - \varphi^{d_{KT}(\sigma_{a_i \leftrightarrow a_{i+1}},\sigma^*)})/Z_\varphi^m = (1 - \varphi) \cdot \varphi^{d_{KT}(\sigma,\sigma^*)}/Z_\varphi^m \\
&= (1 - \varphi) \cdot \Pr[\sigma|\sigma^*].
\end{aligned}$$

The last transition holds trivially (see the proof of Theorem 7.3 for a tighter bound). Thus, we have the desired result. ■ (Theorem 7.6)

While Theorem 7.6 shows that scoring rules such as the Borda count and the harmonic scoring rule have polynomial sample complexity, it does not apply to scoring rules such as plurality and veto since they have $\beta_{\min} = 0$. Note that in Borda count all β_i's are equal, hence it is the rule with the lowest possible $\beta^* = 1$.

7.4 Moving Towards Generalizations

Section 7.3 focused on Mallows' model and sample complexity. In Section 7.5 we will consider a much higher level of abstraction, including much more general noise models and infinitely many samples. This section serves as a mostly conceptual interlude where we gradually introduce some new ideas.

7.4.1 From finite to infinitely many samples and the family of PD-c rules

While the exact or asymptotic sample complexity — as analyzed in Section 7.3 — can help us distinguish between various voting rules, here we take a normative point of view and argue that voting rules need to meet a basic requirement: given *infinitely many* samples, the rule should be able to reproduce the true ranking with probability 1. Formally, a voting rule r is *accurate in the limit* for a noise model G if given votes from G, $\lim_{n\to\infty} \mathrm{Acc}^r(n) = 1$.

For Mallows' model, achieving accuracy-in-the-limit is very easy. Theorem 7.3 shows that given $O(\ln(m/\varepsilon))$ samples, every PM-c rule outputs the true ranking with probability at least $1-\varepsilon$. Thus, every PM-c rule is accurate in the limit for Mallows' model. While plurality requires at least exponentially many samples to determine the true ranking with high probability (Theorem 7.5), a matching upper bound (up to logarithmic factors) can trivially be established showing that plurality is accurate in the limit for Mallows' model as well. In fact, it can be argued that all scoring rules are accurate in the limit for Mallows' model. We prove a more general statement by introducing a novel family of voting rules that generalizes scoring rules and showing that all rules in this family are accurate in the limit for Mallows' model.

Definition 7.3 (Position-Dominance)**.** Given a profile $\pi = (\sigma_1, \ldots, \sigma_n) \in \mathcal{L}(A)^n$, alternative $a \in A$ and $j \in \{1, \ldots, m-1\}$, define $s_j(a) = |\{i : \sigma_i(a) \leqslant j\}|$, i.e., the number of votes in which alternative a is among the first j positions. For $a, b \in A$, we say that *a position-dominates b* if $s_j(a) > s_j(b)$ for all $j \in \{1, \ldots, m-1\}$. The position-dominance graph (PD graph) of π is defined as the directed graph $G = (V, E)$ where alternatives are vertices ($V = A$) and there is an edge from alternative a to alternative b if a position-dominates b.

The concept of position-dominance is reminiscent of the notion of first-order stochastic dominance in probability theory: informally, a random variable (first-order) stochastically dominates another random variable over the same domain if for any value in the domain the former random variable has higher probability of being above the value than the latter random variable. Also note that position-dominance is a transitive relation; for alternatives $a, b, c \in A$ if a position-dominates b and b position-dominates c, then a position-dominates c. However, it is possible that for some alternatives $a, b \in A$, neither a position-dominates b nor b position-dominates a. Thus, the PD graph is always acyclic, but not always complete. When the PD graph is complete, it reduces to a ranking, similarly to the case of the PM graph.

Definition 7.4 (Position-Dominance Consistent Rules)**.** A deterministic voting rule r is called *position-dominance consistent* (PD-c) if $r(\pi) = \sigma$ whenever the PD graph of profile π reduces to ranking σ. For randomized voting rules, we require that $\Pr[r(\pi) = \sigma] = 1$.

This novel family of rules captures voting rules that give higher preference to alternatives that appear at earlier positions. It is quite intuitive that all positional scoring rules are PD-c because they score alternatives purely based on their positions in the rankings and give higher weight to alternatives at earlier positions (a similar observation has been made in [81] in a slightly different context). PD-c rules also capture another classical voting rule — the Bucklin rule. The definition of the Bucklin rule and the proof of Theorem 7.7 appear in Appendix B.3.

Theorem 7.7. *All positional scoring rules and the Bucklin rule are PD-c rules.*

It is easy to argue that all PD-c rules are accurate in the limit for Mallows' model. Let σ^* be the true ranking and a_i be the alternative at position i in σ^*. If we construct a profile by sampling n votes from Mallows' model, then $\mathbb{E}[s_j(a_i)] = n \cdot q_{i,j}$. Recall that $q_{i,j}$ is the probability of alternative a_i appearing among the first j positions in a random vote. Clearly in Mallows' model, $q_{i,j} > q_{l,j}$ for any $i < l$. Therefore, as $n \to \infty$, we will have $\Pr[s_j(a_i) > s_j(a_l)] = 1$ for all $j \in \{1, \ldots, m-1\}$ and $i < l$. Hence, the PD graph of the profile would reduce to σ^* (so any PD-c rule will output σ^*) with probability 1 as $n \to \infty$. We conclude that all PD-c rules are accurate in the limit for Mallows' model.

7.4.2 PM-c rules are disjoint from PD-c rules

In Theorem 7.2 we saw various classical voting rules that are PM-c, and Theorem 7.7 describes well-known voting rules that are PD-c. At first glance, the definitions of PM-c and PD-c may seem unrelated. However, it turns out that no voting rule can be both PM-c and PD-c. To show this we give a carefully constructed profile where both the PM graph and the PD graph are acyclic and complete, but they reduce to different rankings. Hence, a rule that is both PM-c and PD-c must output two different rankings with probability 1, which is impossible. For our example, let $A = \{a, b, c\}$ be the set of alternatives. The profile π consisting of 11 votes is given below.

4 votes	2 votes	3 votes	2 votes
a	b	b	c
b	a	c	a
c	c	a	b

It is easy to check that the PM graph of π reduces to $a \succ b \succ c$ and the PD graph of π reduces to $b \succ a \succ c$. Thus, we have the following result.

Theorem 7.8. *No (randomized) voting rule can be both PM-c and PD-c.*

The theorem is not entirely surprising, as it is known that there is no positional scoring rule that is Condorcet consistent [91]. Note that in addition to PM-c rules and PD-c rules, we can construct numerous simple rules that are also accurate in the limit for Mallows' model, such as the rule that ranks alternatives according to their most frequent position in the input votes and the rule that outputs the most frequent ranking.

7.4.3 Generalizing the noise model

While being accurate in the limit for Mallows' model can be seen as a necessity for voting rules, the assumption that the noise observed in practice would perfectly (or even approximately) fit Mallows' model is unrealistic. For example, Mao et al. [140] show that, in certain real-world scenarios, the noise observed is far from what Mallows predicts. While voting rules cannot be expected to have low sample complexity in all types of noise models that arise in practice, it is reasonable to expect them to be at least accurate in the limit for such noise models.

Unfortunately, it is not clear what noise models can be expected to arise in practice and little attention has been given to characterizing reasonable noise models in the literature. To address this issue we impose a structure, parametrized by distance functions, on the noise models to make them well-behaved. As noted in Section 7.9, this approach is related to the work of Fligner and Verducci [92], but we further generalize the structure of the noise model by removing their assumption of exponentially decreasing probabilities.

Definition 7.5 (d-Monotonic Noise Models)**.** Let σ^* denote the true underlying ranking. Let $d : \mathcal{L}(A) \times \mathcal{L}(A) \to \mathbb{R}_{\geqslant 0}$ be a distance function over rankings. A noise model is called *monotonic* with respect to d (or d-monotonic) if for any $\sigma, \sigma' \in \mathcal{L}(A)$, $d(\sigma, \sigma^*) < d(\sigma', \sigma^*)$ implies $\Pr[\sigma|\sigma^*] > \Pr[\sigma'|\sigma^*]$ and $d(\sigma, \sigma^*) = d(\sigma', \sigma^*)$ implies $\Pr[\sigma|\sigma^*] = \Pr[\sigma'|\sigma^*]$.

In words, given a distance function d we expect that rankings closer to the true ranking would have higher probability of being observed. Note that Mallows' model is monotonic with respect to the KT distance. Fix a distance function d. A noise model that arises in practice can be expected to be monotonic; consequently, we require that a voting rule be accurate in the limit for every d-monotonic noise model.

Definition 7.6. A voting rule r is called *monotone-robust* with respect to distance function d (or d-monotone-robust) if r is accurate in the limit for all d-monotonic noise models.

We saw that all PM-c and PD-c rules are accurate in the limit for Mallows' model. In fact, it can be shown that they are accurate in the limit for all d_{KT}-monotonic noise models, i.e., they are d_{KT}-monotone-robust. However, we omit the proof as the theorem will follow from the even more general results of Section 7.5.

Theorem 7.9. *All PM-c and PD-c rules are d_{KT}-monotone-robust.*

7.5 General Characterizations

For any given distance function d, we proposed d-monotonic noise models in an attempt to capture noise models that may arise in practice. However, until now we only focused on one specific distance function — the KT distance. Noise models parametrized by other distance functions have been studied in the literature starting with Mallows [139] himself. In fact, all our previous proofs relied only on the fact that the KT distance is swap-increasing and Theorem 7.9 can also be shown to hold when the KT distance is replaced by any swap-increasing distance. Alas, among the three most popular distance functions that we consider, only the KT distance is swap-increasing, as shown by the example below.

Example 7.1. Let the set of alternatives be $A = \{a, b, c\}$. Let $\sigma^* = (a \succ b \succ c)$ and $\sigma = (b \succ c \succ a)$. Note that $b \succ_\sigma c$ and $b \succ_{\sigma^*} c$. Now consider the ranking $\sigma_{b \leftrightarrow c} = (c \succ b \succ a)$. It is easy to verify that $d_{FR}(\sigma, \sigma^*) = d_{FR}(\sigma_{b \leftrightarrow c}, \sigma^*) = 4$ and $d_{MD}(\sigma, \sigma^*) = d_{MD}(\sigma_{b \leftrightarrow c}, \sigma^*) = 2$. Thus, the distance does not increase by swapping two alternatives that were in the correct order, which shows that neither the footrule distance nor the maximum displacement distance is swap-increasing.

In this section we ask whether the families of PM-c and PD-c rules are monotone-robust with respect to distance functions other than swap-increasing distances. We fully characterize all distance functions with respect to which all PM-c and/or all PD-c rules are monotone-robust. Given any distance function d, it is easy to construct an equivalent integer-valued distance function d' such that properties like d-monotone-robustness, MC and PC (the latter two are yet to be introduced) are preserved. Thus, without loss of generality we henceforth restrict our distance functions to be integer-valued.

7.5.1 Distances for which all PM-c rules are monotone-robust

We first characterize the distance functions for which all PM-c rules are monotone-robust. This leads us to the definition of a rather natural family of distance functions, which may be of independent interest.

Definition 7.7 (Majority-Concentric (MC) Distances)**.** For any distance function d, ranking $\sigma \in \mathcal{L}(A)$ and integer $k \in \mathbb{N} \cup \{0\}$, let $\mathcal{N}^k(\sigma) = \{\tau \in \mathcal{L}(A) | d(\tau, \sigma) \leqslant k\}$ be the set of all rankings at distance at most k from σ. Furthermore, for any alternatives $a, b \in A$, let $\mathcal{N}^k_{a \succ b}(\sigma) = \{\tau \in \mathcal{N}^k(\sigma) | a \succ_\tau b\}$. A distance function d is called *majority-concentric* (MC) if for any $\sigma \in \mathcal{L}(A)$ and $a, b \in A$ such that $a \succ_\sigma b$, $|\mathcal{N}^k_{a \succ b}(\sigma)| \geqslant |\mathcal{N}^k_{b \succ a}(\sigma)|$ for every $k \in \mathbb{N} \cup \{0\}$.

Consider a ranking σ and imagine *concentric* circles around σ where the kth circle from the center represents the neighborhood $\mathcal{N}^k(\sigma)$. Then, the MC criterion requires that for every pair of alternatives, a (weak) *majority* of rankings in each neighborhood (which can be viewed as a set of votes) agree with σ, hence the name majority-concentric.

There is an alternative and perhaps more intuitive characterization of MC distances. Fix any MC distance d, base ranking σ and alternatives $a, b \in A$ such that $a \succ_\sigma b$. Let $\mathcal{L}_{a \succ b}(A) = \{\tau \in \mathcal{L}(A) | a \succ_\tau b\}$ denote the set of all rankings where $a \succ b$ and let $\mathcal{L}_{b \succ a}(A) = \mathcal{L}(A) \setminus \mathcal{L}_{a \succ b}(A)$. Let us sort all rankings in both sets in the non-decreasing order of their distance from σ, and map the ith ranking (in the sorted order) in $\mathcal{L}_{a \succ b}(A)$ to the ith ranking in $\mathcal{L}_{b \succ a}(A)$. We can show that this mapping takes every ranking to a ranking at equal or greater distance from σ. We call such a mapping *weakly-distance-increasing* with respect to σ. To see this, suppose for contradiction that (say) the ith ranking of $\mathcal{L}_{a \succ b}(A)$ at distance k from σ is mapped to the ith ranking of $\mathcal{L}_{b \succ a}(A)$ at distance $k' < k$ from σ. Then clearly, $|\mathcal{N}^{k'}_{a \succ b}(\sigma)| < i$ and $|\mathcal{N}^{k'}_{b \succ a}(\sigma)| \geqslant i$, which is a contradiction since we assumed the distance to be MC. In the other direction, again fix a distance d, $\sigma \in \mathcal{L}(A)$ and $a, b \in A$ such that $a \succ_\sigma b$. Suppose there exists a bijection $f : \mathcal{L}_{a \succ b}(A) \to \mathcal{L}_{b \succ a}(A)$ that is weakly-distance-increasing with respect to σ. Then for any $k \in \mathbb{N} \cup \{0\}$ we have $\mathcal{N}^k_{b \succ a}(\sigma) \subseteq \{f(\tau) | \tau \in \mathcal{N}^k_{a \succ b}(\sigma)\}$, so $|\mathcal{N}^k_{a \succ b}(\sigma)| \geqslant |\mathcal{N}^k_{b \succ a}(\sigma)|$. If this holds for every $\sigma \in \mathcal{L}(A)$ and $a, b \in A$ such that $a \succ_\sigma b$, then the distance is MC. In conclusion, we have proved the following lemma.

Lemma 7.5. *A distance function d is MC if and only if for every $\sigma \in \mathcal{L}(A)$ and every $a, b \in A$ such that $a \succ_\sigma b$, there exists a bijection $f : \mathcal{L}_{a \succ b}(A) \to \mathcal{L}_{b \succ a}(A)$ that is weakly-distance-increasing with respect to σ.*

We are now ready to prove our first main result of this section: the distance functions with respect to which all PM-c rules are monotone-robust are exactly MC distances.

Theorem 7.10. *All PM-c rules are d-monotone-robust for a distance function d if and only if d is MC.*

Proof. First, we assume that d is MC and show that all PM-c rules are d-monotone-robust. Specifically, consider any d-monotonic noise model G; we wish to show that all PM-c rules are accurate in the limit for G. Let σ^* be an arbitrary true ranking and $a, b \in A$ be two arbitrary alternatives with $a \succ_{\sigma^*} b$.

Using Lemma 7.5, there exists an injection $f : \mathcal{L}_{a \succ b}(A) \to \mathcal{L}_{b \succ a}(A)$ that is weakly-distance-increasing with respect to σ^*. Hence, for every $\sigma \in \mathcal{L}_{a \succ b}(A)$, $d(\sigma, \sigma^*) \leqslant d(f(\sigma), \sigma^*)$, so $\Pr[\sigma|\sigma^*] \geqslant \Pr[f(\sigma)|\sigma^*]$ since G is d-monotonic. Crucially, $\sigma^* \in \mathcal{L}_{a \succ b}(A)$ and $d(\sigma^*, \sigma^*) = 0 < d(f(\sigma^*), \sigma^*)$, so $\Pr[\sigma^*|\sigma^*] > \Pr[f(\sigma^*)|\sigma^*]$. Recall that f is a bijection, hence its range is the whole of $\mathcal{L}_{b \succ a}(A)$. By summing over all $\sigma \in \mathcal{L}_{a \succ b}(A)$, we get

$$\begin{aligned}
\Pr[a \succ b|\sigma^*] = \sum_{\sigma \in \mathcal{L}_{a \succ b}(A)} \Pr[\sigma|\sigma^*] &> \sum_{\sigma \in \mathcal{L}_{a \succ b}(A)} \Pr[f(\sigma)|\sigma^*] \\
&= \sum_{\sigma \in \mathcal{L}_{b \succ a}(A)} \Pr[\sigma|\sigma^*] = \Pr[b \succ a|\sigma^*].
\end{aligned}$$

It follows that given infinitely many samples from G, there would be an edge from a to b in the PM graph with probability 1. Since this holds for all $a, b \in A$, the PM graph would reduce to σ^* with probability 1. Therefore, any PM-c rule would output σ^* with probability 1, as required.

In the other direction, consider any distance function d that is not MC. We show that there exists a PM-c rule that is not accurate in the limit for some d-monotonic noise model G. Since d is not MC, there exists a $\sigma^* \in \mathcal{L}(A)$, an integer k and alternatives $a, b \in A$ with $a \succ_{\sigma^*} b$ such that $|\mathcal{N}^k_{a \succ b}(\sigma^*)| < |\mathcal{N}^k_{b \succ a}(\sigma^*)|$. Now we construct the noise model G as follows. Let $M = \max_{\sigma \in \mathcal{L}(A)} d(\sigma, \sigma^*)$ and let $T > M$ (we will set T later). Define a weight w_σ for each ranking σ as follows: if $d(\sigma, \sigma^*) \leqslant k$ (i.e., $\sigma \in \mathcal{N}^k(\sigma^*)$), then $w_\sigma = T - d(\sigma, \sigma^*)$ else $w_\sigma = M - d(\sigma, \sigma^*)$. Now construct G by assigning probabilities to rankings proportionally to their weights, i.e., $\Pr[\sigma|\sigma^*] = w_\sigma / \sum_{\tau \in \mathcal{L}(A)} w_\tau$. First, by the definition of M and the fact that $T > M$, it is easy to check that G is indeed a probability distribution and that G is d-monotone.

Next, we set T such that $\Pr[a \succ b|\sigma^*] < \Pr[b \succ a|\sigma^*]$. Since the probabilities are proportional to the weights, we want to obtain: $\sum_{\sigma\in\mathcal{L}(A)|a\succ_\sigma b} w_\sigma < \sum_{\sigma\in\mathcal{L}(A)|b\succ_\sigma a} w_\sigma$. Let $|\mathcal{N}^k_{a\succ b}(\sigma^*)| = l$, hence $|\mathcal{N}^k_{b\succ a}(\sigma^*)| \geqslant l+1$. Now, on the one hand,

$$\sum_{\sigma\in\mathcal{L}_{a\succ b}(A)} w_\sigma \leqslant \sum_{\sigma\in\mathcal{N}^k_{a\succ b}(\sigma^*)} T + \sum_{\sigma\in\mathcal{L}_{a\succ b}(A)\setminus\mathcal{N}^k_{a\succ b}(\sigma^*)} M \leqslant l\cdot T + m!\cdot M.$$

On the other hand,

$$\sum_{\sigma\in\mathcal{L}_{b\succ a}(A)} w_\sigma \geqslant \sum_{\sigma\in\mathcal{N}^k_{b\succ a}(\sigma^*)} (T-k) + \sum_{\sigma\in\mathcal{L}_{b\succ a}(A)\setminus\mathcal{N}^k_{b\succ a}(\sigma^*)} 0 \geqslant (l+1)\cdot(T-k).$$

Now we set T such that $(l+1)\cdot(T-k) > l\cdot T + m!\cdot M$, i.e., $T > (l+1)\cdot k + m!\cdot M$. Noting that $l+1 \leqslant m!$ and $k \leqslant M$, we can achieve this by simply setting $T = 2\cdot m!\cdot M$.

Since we have obtained $\Pr[a \succ b|\sigma^*] < \Pr[b \succ a|\sigma^*]$ under G, given infinitely many samples there would be an edge from b to a in the PM graph with probability 1. Therefore, with probability 1 the PM graph would not reduce to σ^*. We can easily construct a PM-c rule r that outputs a ranking σ whenever the PM graph reduces to σ, and outputs an arbitrary ranking with $b \succ a$ when the PM graph does not reduce to any ranking. With probability 1, such a rule would output a ranking where $b \succ a$. Hence, r is not accurate in the limit for G, as required. ■ (Theorem 7.10)

7.5.2 Distances for which all PD-c rules are monotone-robust

We next characterize the distance functions for which all PD-c rules are monotone-robust. This leads us to define another natural family of distance functions.

Definition 7.8 (Position-Concentric (PC) Distances)**.** For any ranking $\sigma \in \mathcal{L}(A)$, integer $k \in \mathbb{N}\cup\{0\}$, integer $j \in \{1,\ldots,m-1\}$ and alternative $a \in A$, let $\mathcal{S}^k_j(\sigma,a) = \{\tau \in \mathcal{N}^k(\sigma)|\tau(a) \leqslant j\}$ be the set of rankings at distance at most k from σ where alternative a is ranked in the first j positions. A distance function d is called *position-concentric* (PC) if for any $\sigma \in \mathcal{L}(A)$, $j \in \{1,\ldots,m-1\}$, and $a,b \in A$ such that $a \succ_\sigma b$, we have that $|\mathcal{S}^k_j(\sigma,a)| \geqslant |\mathcal{S}^k_j(\sigma,b)|$ for all $k \in \mathbb{N}\cup\{0\}$, and strict inequality holds for some $k \in \mathbb{N}\cup\{0\}$.

While MC distances are defined by matching aggregate pairwise comparisons of alternatives in every circle that is centered on the base ranking, PC distances focus on matching pairwise comparisons of aggregate positions of alternatives in every concentric circle. Similarly to Lemma 7.5 for MC distances, PC distances also admit an equivalent characterization. We use this equivalence and show that PC distances are exactly the distance functions with respect to which all PD-c rules are monotone-robust.

Let $\mathcal{S}_j(a) = \{\sigma \in \mathcal{L}(A)|\sigma(a) \leqslant j\}$ denote the set of all rankings where alternative a is ranked among the first j positions. Call a distance function $d : X \to Y$ *distance-increasing* with respect to a ranking σ if $d(f(\tau),\sigma) \geqslant d(\tau,\sigma)$ for every $\tau \in X$ (i.e., d is weakly-distance-increasing) and strict inequality holds for at least one $\tau \in X$.

Lemma 7.6. *A distance function d is PC if and only if for every $\sigma \in \mathcal{L}(A)$, every $a, b \in A$ such that $a \succ_\sigma b$ and every $j \in \{1, \ldots, m-1\}$, there exists a bijection $f : \mathcal{S}_j(a) \to \mathcal{S}_j(b)$ that is distance-increasing with respect to σ.*

Proof. For the forward direction, fix any PC distance d, base ranking σ, alternatives $a, b \in A$ such that $a \succ_\sigma b$ and $j \in \{1, \ldots, m-1\}$. Let us sort all rankings in $\mathcal{S}_j(a)$ and $\mathcal{S}_j(b)$ in the non-decreasing order of their distance from σ, and map the ith ranking (in the sorted order) in $\mathcal{S}_j(a)$ to the ith ranking in $\mathcal{S}_j(b)$. First, we show that this mapping is weakly-distance-increasing with respect to σ. Suppose for contradiction that (say) the ith ranking of $\mathcal{S}_j(a)$ at distance k from σ is mapped to the ith ranking of $\mathcal{S}_j(b)$ at distance $k' < k$ from σ. Then clearly, $|\mathcal{S}_j^{k'}(\sigma, a)| < i$ and $|\mathcal{S}_j^{k'}(\sigma, b)| \geqslant i$, so $|\mathcal{S}_j^{k'}(\sigma, b)| > |\mathcal{S}_j^{k'}(\sigma, a)|$, which is a contradiction since we assumed the distance to be PC. Now we show that this mapping takes at least one ranking to a ranking at strictly greater distance from σ. Since d is PC, there exists some $k^* \in \mathbb{N} \cup \{0\}$ such that $|\mathcal{S}_j^{k^*}(\sigma, a)| > |\mathcal{S}_j^{k^*}(\sigma, b)|$. Consider the largest i such that ith ranking of $\mathcal{S}_j(a)$ is at distance at most k^* from σ. If this ranking is mapped to a ranking at equal distance (at most k^*) from σ, then we would have $|\mathcal{S}_j^{k^*}(\sigma, b)| \geqslant |\mathcal{S}_j^{k^*}(\sigma, a)|$, which is a contradiction. Hence, this ranking is mapped to a ranking at strictly greater distance from σ.

In the other direction, take any distance function d. Suppose for every $\sigma \in \mathcal{L}(A)$, $a, b \in A$ such that $a \succ_\sigma b$ and $j \in \{1, \ldots, m-1\}$, there exists a bijection $f : \mathcal{S}_j(a) \to \mathcal{S}_j(b)$ that is weakly-distance-increasing with respect to σ and maps at least one ranking to a ranking at strictly greater distance from σ. Fix any particular $\sigma \in \mathcal{L}(A)$, $a, b \in A$ such that $a \succ_\sigma b$ and $j \in \{1, \ldots, m-1\}$. First, for any $k \in \mathbb{N}$ we have $\mathcal{S}_j^k(b) \subseteq \{f(\tau) | \tau \in \mathcal{S}_j^k(a)\}$, so $|\mathcal{S}_j^k(a)| \geqslant |\mathcal{S}_j^k(b)|$. Further, take the ranking $\tau \in \mathcal{S}_j(a)$ for which $d(f(\tau), \sigma) > d(\tau, \sigma)$. Let $k^* = d(\tau, \sigma)$. Then, $f(\tau) \in \{f(\gamma) | \gamma \in \mathcal{S}_j^{k^*}(a)\}$ but $f(\tau) \notin \mathcal{S}_j^{k^*}(b)$. Also note that $S_j^{k^*}(b) \subseteq \{f(\gamma) | \gamma \in S_j^{k^*}(a)\}$. Hence, $|\mathcal{S}_j^{k^*}(b)| < |\{f(\gamma) | \gamma \in \mathcal{S}_j^{k^*}(a)\}| = |\mathcal{S}_j^{k^*}(a)|$, as required. Since this holds for every $\sigma \in \mathcal{L}(A)$, $a, b \in A$ such that $a \succ_\sigma b$ and $j \in \{1, \ldots, m-1\}$, the distance is PC. ■ (Lemma 7.6)

We are now ready to prove our characterization result, which is analogous to Theorem 7.10.

Theorem 7.11. *All PD-c rules are d-monotone-robust for a distance function d if and only if d is PC.*

Proof. First, we assume that d is PC and show that all PD-c rules are d-monotone-robust. Consider any d-monotonic noise model G; we wish to show that all PD-c rules are accurate in the limit for G. Fix any true ranking $\sigma^* \in \mathcal{L}(A)$, $a, b \in A$ such that $a \succ_{\sigma^*} b$, and $j \in \{1, \ldots, m-1\}$.

Since d is PC, Lemma 7.6 implies that there exists a bijection $f : \mathcal{S}_j(a) \to \mathcal{S}_j(b)$ such that (i) for every $\sigma \in \mathcal{S}_j(a)$, $d(f(\sigma), \sigma^*) \geqslant d(\sigma, \sigma^*)$, hence $\Pr[\sigma|\sigma^*] \geqslant \Pr[f(\sigma)|\sigma^*]$; and (ii) for some $\sigma \in \mathcal{S}_j(a)$, $d(f(\sigma), \sigma^*) > d(\sigma, \sigma^*)$, hence $\Pr[\sigma|\sigma^*] > \Pr[f(\sigma)|\sigma^*]$. Recall that f is a bijection, hence its range is the whole of $\mathcal{S}_j(b)$. Now we sum over all $\sigma \in \mathcal{S}_j(a)$ (similarly to the proof of Theorem 7.10) and get that the probability that a appears in the first j positions is strictly greater than the probability that b appears in the first j positions in a random vote. It follows that given infinitely many samples from G, a would appear in the first j positions in more votes than b does. Since this holds for all $j \in \{1, \ldots, m-1\}$, there would be an edge from a to b in the PD graph with probability 1. Further, since this holds for all $a, b \in A$, the PD graph would reduce to σ^* with probability 1. Hence, any PD-c rule would output σ^* with probability 1, as required.

In the other direction, consider any distance function d that is not PC. We show that there exists a PD-c rule that is not accurate in the limit for some d-monotonic noise model G. Since d is not PC, there exist $\sigma^* \in \mathcal{L}(A)$, $a, b \in A$ with $a \succ_{\sigma^*} b$ and $j \in \{1, \ldots, m-1\}$ such that either (i) there exists $k^* \in \mathbb{N} \cup \{0\}$ with $|\mathcal{S}_j^{k^*}(\sigma^*, a)| < |\mathcal{S}_j^{k^*}(\sigma^*, b)|$, or (ii) for every $k \in \mathbb{N} \cup \{0\}$, $|\mathcal{S}_j^k(\sigma^*, a)| = |\mathcal{S}_j^k(\sigma^*, b)|$.

In case (i), we construct the noise model G exactly as in the proof of Theorem 7.10. We define $M = \max_{\sigma \in \mathcal{L}(A)} d(\sigma, \sigma^*)$ and $T = 2 \cdot m! \cdot M$. Assign weights $w_\sigma = T - d(\sigma, \sigma^*)$ if $d(\sigma, \sigma^*) \leqslant k^*$ and $w_\sigma = M - d(\sigma, \sigma^*)$ otherwise. The noise model G would consequently assign $\Pr[\sigma|\sigma^*] = w_\sigma / \sum_{\tau \in \mathcal{L}(A)} w_\tau$. It follows that under G, the probability of a appearing in the first j positions of a random vote, i.e., $\sum_{\sigma \in \mathcal{S}_j(a)} \Pr[\sigma|\sigma^*]$, would be strictly less than the probability of b appearing in the first j positions in a random vote, i.e., $\sum_{\sigma \in \mathcal{S}_j(b)} \Pr[\sigma|\sigma^*]$. Thus, given infinitely many samples, b would appear more times in the first j positions than a. This implies that with probability 1, there would be no edge from a to b in the PD graph. Therefore, with probability 1 the PD graph would not reduce to σ^*. We can easily construct a PD-c rule similarly to the proof of Theorem 7.10 that outputs a ranking σ whenever the PD graph reduces to σ, and outputs an arbitrary ranking with $b \succ a$ when the PM graph does not reduce to any ranking. With probability 1, such a rule would output a ranking where $b \succ a$. Hence, r is not accurate in the limit for G, as required.

Consider case (ii). Since $|\mathcal{S}_j^k(\sigma^*, a)| = |\mathcal{S}_j^k(\sigma^*, b)|$ for every $k \in \mathbb{N}$, we also have

$$|\mathcal{S}_j^k(\sigma^*, a)| - |\mathcal{S}_j^{k-1}(\sigma^*, a)| = |\mathcal{S}_j^k(\sigma^*, b)| - |\mathcal{S}_j^{k-1}(\sigma^*, b)|.$$

That is, the number of rankings at distance *exactly* k in which a is in the first j positions is equal to the number of rankings at distance *exactly* k where b is in the first j positions. Now consider any d-monotonic noise model G. Since it assigns equal probabilities to rankings at equal distances, we see that the probability of a appearing in the first j positions of a random vote, i.e., $\sum_{\sigma \in \mathcal{S}_j(a)} \Pr[\sigma|\sigma^*]$, would be exactly equal to the probability of b appearing in the first j positions in a random vote, i.e., $\sum_{\sigma \in \mathcal{S}_j(b)} \Pr[\sigma|\sigma^*]$. Therefore, with probability $1/2$, a would not appear in the first j positions more times than b, in which case there would be no edge from a to b in the PD graph and the PD graph would not reduce to σ^*. Now we can easily construct a PD-c rule r such that r outputs σ whenever the PD graph reduces to σ and outputs a fixed ranking $\sigma' \neq \sigma^*$ whenever the PD graph does not reduce to any ranking. Since the PD graph does not reduce to σ^* with probability at least $1/2$, r is clearly not accurate in the limit under such a noise model. Hence, r is not d-monotone-robust, as required. ■ (Theorem 7.11)

We proved that MC and PC are exactly the distance functions with respect to which all PM-c rules and all PD-c rules, respectively, are monotone-robust. If a distance function d is both MC and PC, then it follows that all PM-c as well as all PD-c rules are d-monotone-robust. On the other hand, if d is not MC (resp., not PC), then there exists a PM-c rule (resp., a PD-c rule) that is not d-monotone-robust. We therefore have the following corollary.

Corollary 7.1. *All rules in the union of PM-c rules and PD-c rules are d-monotone-robust for a distance function d if and only if d is both MC and PC.*

Fix any true ranking $\sigma^* \in \mathcal{L}(A)$ and alternatives $a, b \in A$ such that $a \succ_{\sigma^*} b$. Consider any swap-increasing distance function d. By definition, the mapping which maps every ranking σ with $a \succ_\sigma b$ to the ranking $\sigma_{a\leftrightarrow b}$ increases the distance by at least 1. Therefore it is clearly weakly-distance-increasing with respect to σ^*. Such a mapping is also a bijection from $\mathcal{L}_{a\succ b}(A)$ to $\mathcal{L}_{b\succ a}(A)$. Using Lemma 7.5, it follows that d is MC. While the mapping is also a bijection from $\mathcal{S}_j(a)$ to $\mathcal{S}_j(b)$, it may decrease the distance on $\sigma \in \mathcal{S}_j(a)$ where $b \succ a$. Using additional arguments, however, it is possible to show that d is PC as well. The proof of the following lemma is given in Appendix B.4.

Lemma 7.7. *Any swap-increasing distance function is both MC and PC.*

Corollary 7.1 and Lemma 7.7 imply that all PM-c rules and all PD-c rules are d-monotone-robust for any swap-increasing distance d, which implies Theorem 7.9.

7.5.3 Did we generalize the distance functions enough?

How strong are the characterization results of this section? We saw that all PM-c and PD-c rules are d-monotone-robust for any swap-increasing distance d. However, we remarked at the beginning of this section that we need to widen our family of distances as two of the three popular distances that we study are not swap-increasing. We went ahead and characterized all distance functions for which all PM-c rules or all PD-c rules or both are monotone-robust; respectively, these are all MC distances, all PC distances, and their intersection. Are these families wide enough or do we need to search for better voting rules that work for a bigger family of distance functions? Fortunately, we show that even the intersection of the families of MC and PC distances is sufficiently general to include all three popular distance functions.

Theorem 7.12. *The KT distance, the footrule distance, and the maximum displacement distance are both MC and PC.*

The proof of Theorem 7.12 appears in Appendix B.4. Together with Corollary 7.1, it implies that all PM-c rules and all PD-c rules are monotone-robust with respect to all three popular distance functions that we study. We have established that our new families of distance functions are wide enough; this further justifies our focus on PM-c rules and PD-c rules, as they are monotone-robust with respect to all MC and PC distances, respectively.

That said, it is interesting to consider even wider families of distance functions, which is what we do in the next section.

7.6 Modal Ranking is Unique Within GSRs

In this section, we characterize the modal ranking rule — which selects the most common ranking in a given profile — as the unique rule that is monotone-robust with respect to all distance metrics, among a wide sub-family of GSRs. For this, we use a geometric equivalent of GSRs introduced by Mossel et al. [146] called "hyperplane rules". Like GSRs, hyperplane rules were also originally defined as deterministic SCFs. Below, we give the natural extension of the definition to (possibly randomized) SWFs.

Given a profile π, let x_σ^π denote the fraction of times the ranking $\sigma \in \mathcal{L}(A)$ appears in π. Hence, the point $x^\pi = (x_\sigma^\pi)_{\sigma \in \mathcal{L}(A)}$ lies in a probability simplex $\Delta^{m!}$. This allows us to use rankings from $\mathcal{L}(A)$ to index the $m!$ dimensions of every point in $\Delta^{m!}$. Formally,

$$\Delta^{m!} = \left\{ x \subseteq \mathbb{Q}^{m!} \;\middle|\; \sum_{\sigma \in \mathcal{L}(A)} x_\sigma = 1 \right\}.$$

Importantly, note that $\Delta^{m!}$ contains only points with rational coordinates. Weights $w_\sigma \in \mathbb{R}$ for all $\sigma \in \mathcal{L}(A)$ define a hyperplane H where $H(x) = \sum_{\sigma \in \mathcal{L}(A)} w_\sigma \cdot x_\sigma$ for all $x \in \Delta^{m!}$. This hyperplane divides the simplex into three regions; the set of points on each side of the hyperplane, and the set of points on the hyperplane.

Definition 7.9 (Hyperplane Rules)**.** *A hyperplane rule is given by $r = (\mathcal{H}, g)$, where $\mathcal{H} = \{H_i\}_{i=1}^{l}$ is a finite set of hyperplanes, and $g : \{+, 0, -\}^l \to \mathcal{D}(\mathcal{L}(A))$ is a function that takes as input the signs of all the hyperplanes at a point and returns a distribution over rankings. Thus, $r(\pi) = g(\mathtt{sgn}(\mathcal{H}(x^\pi)))$, where*

$$\mathtt{sgn}(\mathcal{H}(x^\pi)) = (\mathtt{sgn}(H_1(x^\pi)), \ldots, \mathtt{sgn}(H_l(x^\pi))),$$

and $\mathtt{sgn} : \mathbb{R} \to \{+, 0, -\}$ is the signum function given by

$$\mathtt{sgn}(x) = \begin{cases} + & x > 0 \\ 0 & x = 0 \\ - & x < 0 \end{cases}$$

Next, we state the equivalence between hyperplane rules and GSRs in the case of randomized SWFs. This equivalence was established by Mossel et al. [146] for deterministic SCFs; it uses the output of a given GSR for each set of compatible vectors to construct the output of its corresponding hyperplane rule in each region, and vice-versa. Simply changing the output of the g functions of both the GSR and the hyperplane rule from a winning alternative (for deterministic SCFs) to a distribution over rankings (for randomized SWFs) and keeping the rest of the proof intact shows the equivalence for randomized SWFs.

Lemma 7.8. *[146] For randomized social welfare functions, the class of generalized scoring rules coincides with the class of hyperplane rules.*

We impose a technical restriction on GSRs that has a clear interpretation under the geometric hyperplane equivalence. Intuitively, it states that if the rule outputs the same ranking (without ties) almost everywhere around a point x^π in the simplex, then the rule must output the same ranking (without ties) on π as well. More formally, consider the regions in which the simplex is divided by a set of hyperplanes $\mathcal{H}$. We say that a region is *interior* if none of its points lie on any of the hyperplanes in $\mathcal{H}$, that is, if for every point x in the region, $\texttt{sgn}(\mathcal{H}(x))$ does not contain any zeros. For $x \in \Delta^{m!}$, let

$$\mathcal{S}(x) = \left\{ y \in \Delta^{m!} \;\middle|\; \forall \sigma \in \mathcal{L}(A),\ x_\sigma = 0 \Rightarrow y_\sigma = 0 \right\}$$

denote the subspace of points that are zero in every coordinate where x is zero. We say that an interior region is *adjacent* to x if its intersection with $\mathcal{S}(x)$ contains points arbitrarily close to x.

Definition 7.10 (No Holes Property)**.** *We say that a hyperplane rule (generalized scoring rule) has no holes if it outputs a ranking σ with probability* 1 *on a profile π whenever it outputs σ with probability* 1 *in all interior regions adjacent to x^π.*

When this property is violated, we have a point x^π such that the output of the rule on x^π is different from the output of the rule almost everywhere around x^π, creating a hole at x^π. We later show (in Section 7.7; see Theorem 7.14) that the no holes property is a very mild restriction on GSRs. We are now ready to formally state our main result.

Theorem 7.13. *Let r be a (possibly randomized) generalized scoring rule with no holes. Then, r is monotone-robust with respect to all distance metrics if and only if r coincides with the modal ranking rule on every profile with no ties (i.e., r outputs the most frequent ranking with probability* 1 *on every profile where it is unique).*

Before proving the theorem, we wish to point out three subtleties. First, our assumption of accuracy in the limit imposes a condition on the rule as the number of votes *goes to infinity*. This has to be translated into a condition on all *finite profiles*; we do this by leveraging the structure of generalized scoring rules.

Second, if there are several rankings that appear the same number of times, a monotone-robust rule can actually output any ranking with impunity, because in the limit this event happens with probability zero.

Third, every noise model G that is monotone with respect to some distance metric satisfies $\Pr_G[\sigma^*;\sigma^*] > \Pr_G[\sigma;\sigma^*]$ for all pairs of different rankings $\sigma, \sigma^* \in \mathcal{L}(A)$. It seems intuitive that the converse holds, i.e., if a noise model satisfies $\Pr_G[\sigma^*;\sigma^*] > \Pr_G[\sigma;\sigma^*]$ for all $\sigma \neq \sigma^*$ then there exists a distance metric d such that G is monotone with respect to d — but this is false. Hence, our condition asks for accuracy in the limit for noise models that are monotone with respect to some metric, instead of just assuming accuracy in the limit with respect to *all* noise models where the ground truth is the unique mode.

Proof of Theorem 7.13. Let r be a (possibly) randomized generalized scoring rule with no holes. Using Lemma 7.8, we represent r as a hyperplane rule $(\mathcal{H}, g)$ where $\mathcal{H} = \{H_i\}_{i=1}^{l}$ is the set of hyperplanes.

First, we show the simpler forward direction. Let r output the most frequent ranking with probability 1 on every profile where it is unique. We want to show that r is monotone-robust with respect to all distance metrics. Take a distance metric d, a d-monotonic noise model G, and a true ranking σ^*. We need to show that r outputs σ^* with probability 1 given infinitely many samples from $G(\sigma^*)$.

Note that d satisfies $d(\sigma^*, \sigma^*) = 0 < d(\sigma, \sigma^*)$ for all $\sigma \neq \sigma^*$. Hence, G must satisfy $\Pr_G[\sigma^*; \sigma^*] > \Pr_G[\sigma; \sigma^*]$ for all $\sigma \neq \sigma^*$. Now, given infinite samples from $G(\sigma^*)$, σ^* becomes the unique most frequent ranking with probability 1. Thus, r outputs σ^* with probability 1 in the limit, as required.

For the reverse direction, let r be d-monotone-robust for all distance metrics d. Take a profile π^* with a unique most frequent ranking σ^*. Recall that $x_\sigma^{\pi^*}$ denotes the fraction of times σ appears in π^* and note that $x_{\sigma^*}^{\pi^*} > x_\sigma^{\pi^*}$ for all $\sigma \neq \sigma^*$. We also denote by $X_\sigma^{\pi^*}$ the number of times σ appears in π^*.

The rest of the proof is organized in three steps. First, we define a distance metric d, a d-monotonic noise model G, and a true ranking. Second, we show that given samples from $G(\sigma^*)$, in the limit r outputs σ^* with probability 1 in every interior region adjacent to x^{π^*}. Finally, we use the no holes property of r to argue that $\Pr[r(\pi^*) = \sigma^*] = 1$.

Step 1: We define d as

$$d(\sigma, \sigma') = \begin{cases} \max(1, |X_\sigma^{\pi^*} - X_{\sigma'}^{\pi^*}|) & \text{if } \sigma \neq \sigma', \\ 0 & \text{otherwise.} \end{cases}$$

We claim that d is a distance metric. Indeed, the first two axioms are easy to verify. The triangle inequality $d(\sigma, \sigma') \leqslant d(\sigma, \sigma'') + d(\sigma'', \sigma')$ holds trivially if any two of the three rankings are equal. When all three rankings are distinct,

$$\begin{aligned} d(\sigma, \sigma'') + d(\sigma'', \sigma') &= \max(1, |X_\sigma^{\pi^*} - X_{\sigma''}^{\pi^*}|) + \max(1, |X_{\sigma''}^{\pi^*} - X_{\sigma'}^{\pi^*}|) \\ &\geqslant \max(1 + 1, |X_\sigma^{\pi^*} - X_{\sigma''}^{\pi^*}| + |X_{\sigma''}^{\pi^*} - X_{\sigma'}^{\pi^*}|) \\ &\geqslant \max(1, |X_\sigma^{\pi^*} - X_{\sigma'}^{\pi^*}|) = d(\sigma, \sigma'). \end{aligned}$$

Now, define the noise model G where

$$\Pr_G[\sigma; \sigma'] = \frac{1/(1 + d(\sigma, \sigma'))}{\sum_{\tau \in \mathcal{L}(A)} 1/(1 + d(\tau, \sigma'))} \quad \text{for } \sigma' \neq \sigma^*.$$

and $\Pr_G[\sigma; \sigma^*] = x_\sigma^{\pi^*}$. Note that G is trivially d-monotone for true rankings other than σ^*. Denoting the number of votes in π^* by n^*, since σ^* is the unique most frequent ranking, we have that $d(\sigma, \sigma^*) = n^*(x_{\sigma^*}^{\pi^*} - x_\sigma^{\pi^*})$ for all $\sigma \neq \sigma^*$. Hence, $\Pr_G[\sigma_1; \sigma^*] \geqslant \Pr_G[\sigma_2; \sigma^*]$ if and only if $d(\sigma_1, \sigma^*) \leqslant d(\sigma_2, \sigma^*)$ and G is also d-monotone for the true ranking σ^*. We conclude that G is a d-monotonic noise model.

Step 2: Let π_n denote a profile consisting of n i.i.d. samples from $G(\sigma^*)$. Since r is monotone-robust for every distance metric, we have

$$\lim_{n \to \infty} \Pr[r(\pi_n) = \sigma^*] = 1. \tag{7.4}$$

If π^* has only one ranking, then only that ranking will ever be sampled. Hence, we will have $\Pr[x^{\pi_n} = x^{\pi^*}] = 1$, and Equation (7.4) would imply that the rule must output σ^* with probability 1 on π^*.

Assume that π^* has at least two distinct votes. We want to show that r outputs σ^* with probability 1 in every interior region adjacent to x^{π^*}. As $n \to \infty$, the distribution of x^{π_n} tends to a Gaussian with mean x^{π^*} and concentrated on the hyperplane

$$\sum_{\sigma \in \mathcal{L}(A) | x_\sigma^{\pi^*} > 0} x_\sigma^{\pi_n} = 1.$$

This follows from the multivariate central limit theorem; see [146] for a detailed explanation. Note that the sum ranges only over the rankings that appear in π^* because in the distribution $G(\sigma^*)$, the probability of sampling a ranking σ that does not appear in π^* is zero.

Since the Gaussian lies in the subspace $\mathcal{S}(x^{\pi^*})$, we set the coordinates corresponding to rankings that do not appear in π^* to zero in all the hyperplanes, and remove the hyperplanes that become trivial. Hereinafter we only consider the rest of the hyperplanes, and the regions they form around x^{π^*}, all in the subspace $\mathcal{S}(x^{\pi^*})$.

If none of the hyperplanes pass through x^{π^*}, then there is a unique interior region K which actually contains x^{π^*} as its interior point. In this case, the limiting probability of x^{π_n} falling in K will be 1, as the Gaussian becomes concentrated around x^{π^*}. Thus, Equation (7.4) implies that r outputs σ^* with probability 1 in K, and therefore on π^*.

If there exists a hyperplane passing through x^{π^*}, then each interior region K adjacent to x^{π^*} is the intersection of finitely many halfspaces whose hyperplanes pass through x^{π^*}. Let $\overline{K}$ and $\overline{\mathcal{S}(x^{\pi^*})}$ denote the closures of K and $\mathcal{S}(x^{\pi^*})$ in $\mathbb{R}^{m!}$, respectively.[6] Thus, $\overline{K}$ is a pointed convex cone with its apex at x^{π^*}, and must subtend a positive solid angle (in $\overline{\mathcal{S}(x^{\pi^*})}$) at its apex since the hyperplanes are distinct. By definition, the solid angle that $\overline{K}$ forms at x^{π^*} is the fraction of volume (the Lebesgue measure in $\overline{\mathcal{S}(x^{\pi^*})}$) covered by $\overline{K}$ in a ball of radius ρ (again in $\overline{\mathcal{S}(x^{\pi^*})}$) centered at x^{π^*}, as $\rho \to 0$ [see, e.g., Section 2 in 74].

Since the Gaussian is symmetric in $\overline{\mathcal{S}(x^{\pi^*})}$ around x^{π^*}, and the limiting distribution of x^{π_n} converges to the Gaussian, the limiting probability of x^{π_n} lying in K is positive. This holds for every interior region K adjacent to x^{π^*}. Thus, Equation (7.4) again implies that r outputs σ^* with probability 1 in every interior region adjacent to x^{π^*}.

Step 3: Finally, since r has no holes and it outputs σ^* with probability 1 in every interior region adjacent to x^{π^*}, we conclude that r must also output σ^* with probability 1 on π^*. ■(Proof of Theorem 7.13)

[6] We remark that considering the closures is necessary since $\Delta^{m!}$ contains only points with rational coordinates; hence it (as well as any subset of it) has measure zero.

7.7 How Restrictive is the No Holes Property?

To complete the picture, we wish to show that the no holes condition that Theorem 7.13 imposes on GSRs is indeed unrestrictive, by establishing that many prominent voting rules (in the sense of receiving attention in the computational social choice literature) are GSRs with no holes. One issue that must be formally addressed is that the definitions of prominent voting rules (see Section 7.2) typically do not address how ties are broken. For example, the plurality rule ranks the alternatives by their number of voters who rank them first; but what should we do in case of a tie? A common practice is to adopt uniformly random tie-breaking; this is almost always used in political elections (e.g., by throwing dice or drawing cards in small municipal elections where ties are not unlikely to occur). From a theoretical point of view, randomized tie-breaking is necessary in order to achieve neutrality with respect to the alternatives [149].

We show that prominent voting rules are GSRs with no holes under a wide family of randomized tie-breaking schemes, which we call *inclusive tie-breaking*.

Definition 7.11 (Inclusive tie-breaking)**.** A tie-breaking scheme is called *inclusive* if it assigns a positive probability to each of the tied decisions at every stage.

Such tied decisions could vary for different rules. For rules that assign scores to alternatives and order them according to their scores, the decisions correspond to choosing the order of alternatives with identical scores. For rules that assign scores to rankings and choose the ranking with the highest score, the decisions correspond to breaking ties among rankings with identical highest scores. While most voting rules use tie-breaking only once (either initially or in the end), multi-stage protocols such as single transferable vote (STV) use tie-breaking in each stage.

We note that uniformly random tie-breaking, which assigns equal probability to every decision, is a special case of inclusive tie-breaking. Admittedly, inclusive tie-breaking does not capture deterministic tie-breaking schemes (such as lexicographic tie-breaking). However, we strongly believe that prominent voting rules other than the modal ranking rule are not monotone-robust with respect to all distance metrics even if a deterministic tie-breaking scheme were used.

Before we demonstrate that prominent voting rules are GSRs with no holes, we show that the no holes condition is implied by a property well-known in the social choice literature as *consistency*. This yields a way of leveraging known results from the literature to easily establish that all positional scoring rules, the Kemeny rule, and single transferable vote (STV) are GSRs with no holes. Intuitively, consistency means that if the output of a rule is identical on two profiles, then taking their union should not change the output. For deterministic SWFs, consistency was first studied by Young and Levenglick [210], who observed that it is incomparable to, but usually much weaker than, consistency of winning alternative in the case of SCFs. Later, Conitzer and Sandholm [66] showed that consistency (whether in rankings or in winning alternatives) is a necessary condition for a voting rule to be a maximum likelihood estimator under i.i.d. votes. We formalize a related, but weaker notion of consistency for randomized SWFs.

Definition 7.12 (Weak consistency for rankings)**.** A randomized SWF r is said to satisfy weak consistency for rankings if $\Pr[r(\pi_1) = \sigma] = 1$ and $\Pr[r(\pi_2) = \sigma] = 1$ implies $\Pr[r(\pi_1 \cup \pi_2) = \sigma] = 1$ for all profiles π_1 and π_2, and all rankings $\sigma \in \mathcal{L}(A)$. Here, $\pi_1 \cup \pi_2$ denotes the profile representing the union of π_1 and π_2.

For hyperplane rules (generalized scoring rules), weak consistency for rankings is equivalent to convexity of the region where the rule outputs σ with probability 1, for every $\sigma \in \mathcal{L}(A)$. Now, we are ready to prove the following implication.

Lemma 7.9. *Any generalized scoring rule satisfying weak consistency for rankings has no holes.*

Proof. Take a GSR r that satisfies weak consistency for rankings. Suppose for contradiction that r has a hole at x^π for some profile π. Let r output σ with probability 1 in all interior regions adjacent to x^π, but not on π. Let k be the number of distinct rankings that appear in π. Hence, $\mathcal{S}(x^\pi)$ is a k-dimensional subspace of $\Delta^{m!}$.

If $k = 1$, then π has only one distinct ranking σ. Thus, $x^\pi_\sigma = 1$ and $x^\pi_{\sigma'} = 0$ for all $\sigma' \neq \sigma$. By the definition of $\mathcal{S}(x^\pi)$, for every $y \in \mathcal{S}(x^\pi)$ we have $y_{\sigma'} = 0$ for all $\sigma' \neq \sigma$. Thus, $y_\sigma = 1$, implying that $\mathcal{S}(x^\pi) = \{x^\pi\}$. Hence, trivially, x^π cannot be a hole.

Let $k > 1$. Define

$$V = \left\{ v \in \{-1, 0, 1\}^{m!} \;\middle|\; \forall \sigma \in \mathcal{L}(A),\ v_\sigma = 0 \Leftrightarrow x^\pi_\sigma = 0 \right.$$
$$\left. \wedge\, \exists \sigma \in \mathcal{L}(A),\ v_\sigma = 1 \;\wedge\; \exists \sigma \in \mathcal{L}(A),\ v_\sigma = -1 \right\}.$$

Now, for every $v \in V$, define the orthant

$$O^v = \left\{ y \in \mathcal{S}(x^\pi) \;\middle|\; \forall \sigma \in \mathcal{L}(A),\ (v_\sigma = 1 \Rightarrow y_\sigma > x^\pi_\sigma) \wedge (v_\sigma = -1 \Rightarrow y_\sigma < x^\pi_\sigma) \right\}.$$

Note that we do not consider the orthants where all the k coordinates are higher (respectively, lower) than those of x^π because such orthants do not have any points in $\mathcal{S}(x^\pi)$ as the sum of those k coordinates must be equal to 1. The rest $3^k - 2$ orthants have non-empty intersection with $\mathcal{S}(x^\pi)$. Further, since the interior regions adjacent to x^π surround it in the space $\mathcal{S}(x^\pi)$ and so do the $3^k - 2$ orthants, each orthant O^v must have a point x^v in some interior region adjacent to x^π, where the output is σ. Now, a convexity lemma (Lemma B.1 in the appendix) shows that we can get x^π as a convex combination of points in $\{x^v | v \in V\}$,[7] on all of which r outputs σ with probability 1. Hence, due to weak consistency for rankings, r must also output σ with probability 1 on x^π, a contradiction. Hence, r has no holes. ■(Proof of Lemma 7.9)

Finally, we are ready to prove that prominent voting rules are GSRs with no holes under all inclusive tie-breaking schemes (which contain uniformly random tie-breaking).

Theorem 7.14. *Under any inclusive tie-breaking scheme, all positional scoring rules, the Kemeny rule, STV, Copeland's method, Bucklin's rule, the maximin rule, Slater's rule, and the ranked pairs method are generalized scoring rules with no holes.*

[7]In Lemma B.1, take $FIX = \{\sigma \in \mathcal{L}(A) | x^\pi_\sigma = 0\}$.

Proof. It can easily be checked that the f functions of the GSR constructions given by Xia and Conitzer [204] and the hyperplanes for the hyperplane rule constructions given by Mossel et al. [146] encode enough information (including all the ties) in their input to the g functions such that it is possible to change the output of g from a winning alternative (for deterministic SCFs) to any desired distribution over rankings (for randomized SWFs) for the rules mentioned in the statement of the theorem. In particular, these rules can be implemented with any inclusive tie-breaking scheme as GSRs.

It is well-known and easy to check that all positional scoring rules are consistent for rankings [see 66, 70]. Young and Levenglick [210] showed that the Kemeny rule is also consistent for rankings. Conitzer and Sandholm [66] showed that STV is consistent for rankings; later it was shown that STV is consistent for rankings only in the absence of tie-breaking, if another mild condition called *continuity* is required [70]. It can be checked that under inclusive tie-breaking, STV returns a single ranking with probability 1 if and only if there are no ties. Hence, consistency in the absence of tie-breaking implies weak consistency for rankings. Thus, we have the following.

Lemma 7.10. *Under any inclusive tie-breaking scheme, all positional scoring rules, the Kemeny rule, and single transferable vote (STV) satisfy weak consistency for rankings.*

Hence, the no holes property of all positional scoring rules, the Kemeny rule, and single transferable vote (STV) follows by Lemmas 7.9 and 7.10. Conitzer and Sandholm [66] showed that other rules such as Bucklin's rule, Copeland's method, the maximin rule, and the ranked pairs method are not consistent for rankings even in the absence of ties. Hence, these rules do not satisfy weak consistency for rankings. Still, we will show that they satisfy the no holes property as well, albeit using a different (and significantly more involved) approach.

Take a hyperplane rule r. We want to show that r does not have a hole at a profile π. Let $k = |\{\sigma \in \mathcal{L}(A) | x^{\pi}_{\sigma} > 0\}|$ be the number of distinct rankings in π. If $k = 1$, then as shown in the proof of Lemma 7.9, $\mathcal{S}(x^{\pi}) = \{x^{\pi}\}$, and there cannot be a hole at π. Assume $k \geqslant 2$. For a set of profiles P, let $x^P = \{x^{\pi'} | \pi' \in P\}$.

Let $dim(\cdot)$ denote the Hausdorff dimension of a given subset of $\mathbb{R}^{m!}$. For any set $C \subseteq \Delta^{m!}$, let $\overline{C}$ denote its closure in $\mathbb{R}^{m!}$. Hence, we have that $dim(\overline{\mathcal{S}(x^{\pi})}) = k - 1$, because $\mathcal{S}(x^{\pi})$ has $k - 1$ free variables.

Lemma 7.11. *Let T denote the set of points in $\mathcal{S}(x^{\pi})$ that lie on at least one of the hyperplanes of r. Then, $dim(\overline{T}) \leqslant k - 2$.*

Proof. Take a hyperplane $\sum_{\sigma \in \mathcal{L}(A)} w_{\sigma} x_{\sigma} = 0$ of r. Consider its intersection with $\overline{\mathcal{S}(x^{\pi})}$. First, we notice that all but k of the x_{σ}'s must be set to zero. Among the remaining k, if we substitute values for $k - 2$ of the variables, we get two equations in two variables, which can be seen to be independent since the one obtained from the hyperplane has the RHS zero, while the one obtained from $\overline{\mathcal{S}(x^{\pi})}$ has the RHS one. Hence, there is at most one solution of the pair of equations.

That is, every combination of values of $k - 2$ free variables lead to at most one solution for the remaining variables. Thus, the dimension of the intersection of the hyperplane with $\overline{\mathcal{S}(x^{\pi})}$ is at most $k - 2$. Taking union over finitely many hyperplanes does not increase the Hausdorff dimension. Hence, we have $dim(\overline{T}) \leqslant k-2$. ■(Proof of Lemma 7.11)

Next, we describe an outline that we follow in order to prove that the no holes property is satisfied by many prominent voting rules. We consider rules that assign a score to every alternative, and then order them in a non-increasing or non-decreasing order of their scores, breaking ties among alternatives with identical scores. This applies to Copeland's method, the maximin rule, and Bucklin's rule.[8] Such rules return a single ranking with probability 1 if and only if the scores of the alternatives are strictly ordered according to that ranking. Let us denote the score of alternative c in profile π by $SC^{\pi}(c)$.

1. For the sake of contradiction, we assume that the rule under consideration, say r, has a hole at a profile π. Hence, r outputs a ranking τ with probability 1 in every interior region adjacent to x^{π}, but there exists a ranking $\tau' \neq \tau$ such that $\Pr[r(\pi) = \tau'] > 0$.
2. Since $\tau' \neq \tau$, there must exist alternatives a and b such that $a \succ_{\tau} b$, but $b \succ_{\tau'} a$. Due to the inclusive tie-breaking scheme, we must have that
   - $SC^{\pi}(b) \geqslant SC^{\pi}(a)$, and
   - $SC(a) > SC(b)$ in every interior region adjacent to x^{π}.
3. Finally, we find a neighborhood of π where we also have $SC(b) \geqslant SC(a)$. Formally, we find a set of profiles P such that
   - for every profile $\pi' \in P$, $SC^{\pi'}(b) \geqslant SC^{\pi'}(a)$, and $x^{\pi'}$ either lies in an interior region adjacent to x^{π} or on one of the hyperplanes of r, and
   - $dim(\overline{x^P}) = k-1$.

Given this, we argue that a contradiction can be reached. Recall that T is the intersection of the hyperplanes of r with $\mathcal{S}(x^{\pi})$. Suppose that $x^P \subseteq T$. Then, $\overline{x^P} \subseteq \overline{T}$, which is impossible because $dim(\overline{x^P}) > dim(\overline{T})$ (Lemma 7.11). Hence, there must exist a profile π' such that $x^{\pi'} \in x^P \setminus T$ lies in an interior region adjacent to x^{π}. However $SC^{\pi'}(b) \geqslant SC^{\pi'}(a)$, which is the desired contradiction.

Note that the first two steps are common to all voting rules. All we need to do is to find a set of profiles P satisfying the stated conditions. For many of the voting rules, P is obtained by increasing $x^{\pi}_{\sigma^*}$ for some $\sigma^* \in \pi$, and decreasing x^{π}_{σ} for all $\sigma \neq \sigma^*$ that appear in π. Formally,

$$P = \left\{ \pi' \;\middle|\; \forall \sigma \in \mathcal{L}(A), \right.$$

$$\left. x^{\pi'}_{\sigma} = \begin{cases} 0, & \text{if } x^{\pi}_{\sigma} = 0, \\ x^{\pi}_{\sigma^*} + \delta, & \text{if } \sigma = \sigma^*, \\ x^{\pi}_{\sigma} - \delta_{\sigma}, & \text{otherwise}, \end{cases} \quad \text{where } 0 < \delta \leqslant \delta_{\max} \wedge \sum_{\sigma \neq \sigma^*, x^{\pi}_{\sigma} > 0} \delta_{\sigma} = \delta \right\},$$

[8]We will see that Slater's rule, which assigns a score to every ranking, can also be handled in this outline.

By the construction, for every profile $\pi' \in P$, the weights of the edges of the weighted PM graph of π increase in the direction of σ^*, and decrease in the direction opposite to σ^* (except the edges with weights 1 and 0 do not change). If $\delta_{\max}$ is chosen to be small enough, this change does not alter the direction of any existing edge in the unweighted PM graph, but breaks all existing ties between pairs of alternatives in one direction or the other. Clearly, $dim(\overline{x^P}) = k - 1$ since decreasing the fractions of all rankings $\sigma \neq \sigma^*$ with $x^{\pi}_{\sigma} > 0$ so that the decrements sum up to δ gives $k - 2$ degrees of freedom choosing δ gives another degree of freedom. This observation is very crucial to the proofs for many of the voting rules.

Below, we describe appropriate choices of σ^* and $\delta_{\max}$ for various prominent voting rules, namely for Copeland's method, Bucklin's rule, the maximin rule, and Slater's rule. We also provide a proof for the ranked pairs method in which we use completely different arguments.

(a) Copeland's method.

Recall that for Copeland's method, $SC^{\pi}(c)$ is the number of outgoing edges from c in the unweighted PM graph of π. If there are no ties in the unweighted PM graph of π, then choosing any $\sigma^* \in \pi$ and a small enough $\delta_{\max}$ ensures that the set P obtained fits the requirements of step 3 of the outline and preserves all the edges in the unweighted PM graph. Hence, $SC(b) \geqslant SC(a)$ is preserved, as required.

In case of ties, let $TIE^{\pi}(c)$ be the set of alternatives with which c is tied in the unweighted PM graph of π.[9] For $\sigma \in \mathcal{L}(A)$, let

$$s(\sigma) = \sum_{c \in TIE^{\pi}(b)} \mathbb{I}[b \succ_{\sigma} c] - \sum_{c \in TIE^{\pi}(b)} \mathbb{I}[c \succ_{\sigma} b] - \sum_{c \in TIE^{\pi}(a)} \mathbb{I}[a \succ_{\sigma} c] + \sum_{c \in TIE^{\pi}(a)} \mathbb{I}[c \succ_{\sigma} a].$$

Let $n^{\pi}_{x \succ y}$ denote the number of rankings that prefer alternative x to alternative y in π. Summing over all rankings in π and changing the order of the summation in each term, we get

$$\sum_{i=1}^{n} s(\sigma_i) = \sum_{c \in TIE^{\pi}(b)} \left(n^{\pi}_{b \succ c} - n^{\pi}_{c \succ b} \right) - \sum_{c \in TIE^{\pi}(a)} \left(n^{\pi}_{a \succ c} - n^{\pi}_{c \succ a} \right) = 0,$$

where the last equality holds due to the definitions of $TIE^{\pi}(b)$ and $TIE^{\pi}(a)$. Also, note that the sum evaluates to zero even if either $TIE^{\pi}(b)$ or $TIE^{\pi}(a)$ or both are empty sets.

Hence, there exists a ranking $\sigma^* \in \pi$ such that $s(\sigma^*) \geqslant 0$. There exists a $\delta_{\max} > 0$ such that increasing $x^{\pi}_{\sigma^*}$ by at most $\delta_{\max}$ and decreasing the fractions of other rankings that appear in π would not change the non-tied edges of the PM graph, and among the ties, b would defeat at least as many previously tied alternatives as a does. Hence, such a change preserves $SC(b) \geqslant SC(a)$. Further, $\delta_{\max}$ is chosen to be small enough so that for the new profile π', $x^{\pi'}$ does not fall in an interior region that is not adjacent to x^{π}, i.e., it either lies in an interior region adjacent to x^{π} or on one of the hyperplanes of Copeland's method. Thus, the set of profiles P obtained in this way fits the requirements of step 3 of the outline.

[9] We add zero to the Copeland score of an alternative for its tied edges; this is also known as $Copeland^0$.

(b) Bucklin's rule.

Let $SC^{\pi}(a) = k$. We know that $SC^{\pi}(b) \leqslant SC^{\pi}(a) = k$.[10] Let $T^{\pi}(j, c)$ denote the fraction of rankings where c is ranked in the first j positions. Then, by the definition of the Bucklin score,

$$T^{\pi}(k, b) > 1/2 \quad \text{and} \quad T^{\pi}(k-1, a) \leqslant 1/2. \tag{7.5}$$

If we find σ^* such that the set P defined in the outline preserves the two inequalities in Equation (7.5), then we will have $SC(a) \geqslant k$ and $SC(b) \leqslant k$, i.e., $SC(b) \leqslant SC(a)$ will be preserved.

Let $T^{\pi}(k, b) = 1/2 + \gamma$. Then, it is easy to check that if the fractions of all the rankings in π are altered by less than $\gamma/m!$, then we would still have $T(k, b) > 1/2$. Now, we simply observe that since $T^{\pi}(k-1, a) \leqslant 1/2$, more than half of the rankings in π, in particular, at least one ranking ranks a not in the first $k-1$ positions. Choosing this as σ^* and taking $\delta_{\max} < \gamma/m!$ (and also small enough so that the new profile does not lie in an interior region not adjacent to x^{π}) would preserve both inequalities in Equation (7.5).

(c) The maximin rule.

Here, $SC^{\pi}(c)$ is the minimum of the weights of the outgoing edges from c in the weighted PM graph of π. Let $MINW^{\pi}(c)$ denote the set of alternatives to which c has an outgoing edge with the minimum weight in the weighted PM graph of π. Now, take an alternative $c \in MINW^{\pi}(a)$. Let w be the weight of the edge from a to c. First, we note that $w \neq 1$, because $w = 1$ would imply that a has an outgoing edge with weight 1 to every other alternative, i.e., a is ranked first in all votes in π. This would contradict $SC^{\pi}(b) \geqslant SC^{\pi}(a)$. Next, if $w = 0$, then c beats a in every vote in π. Now, all profiles in $\mathcal{S}(x^{\pi})$ have the same set of rankings as π, and hence have zero maximin score of a. Thus, $SC(b) \geqslant SC(a)$ is trivially satisfied in any point of $\mathcal{S}(x^{\pi})$ and, subsequently, we can define P so that x^P is the union of the interior regions adjacent to x^{π}.

Let us assume $w \in (0, 1)$. Let $R_{c \succ a}(\pi)$ be the set of rankings in π where $c \succ a$, and define $R_{a \succ c}$ to be the set of rankings in π where $a \succ c$. Since $w \in (0, 1)$, $R_{a \succ c} \neq \emptyset$ and $R_{c \succ a} \neq \emptyset$. To obtain P, we do not choose one $\sigma^* \in \pi$, increase its fraction and decrease the fractions of the rest of the rankings in π. Rather, we increase the fractions of all rankings in $R_{c \succ a}$ by a total of δ, and decrease the fractions of all rankings in $R_{a \succ c}$ by a total of δ, where $0 < \delta \leqslant \delta_{\max}$. Once again, we choose $\delta_{\max} > 0$ small enough so that x^P does not intersect with interior regions not adjacent to x^{π}. Increasing the fractions of all rankings $R_{c \succ a}$ so that the increments add up to δ gives $|R_{c \succ a}| - 1$ degrees of freedom. Similarly, decreasing the fractions of all rankings in $R_{a \succ c}$ so that the decrements add up to δ gives another $|R_{a \succ c}| - 1$ degrees of freedom. Finally, choosing δ itself gives one degree of freedom. Hence, the set of profiles P obtained satisfy $dim(\overline{x^P}) = k - 1$.

Further, note that by construction, the weight of the edge from a to c drops by δ. Hence, the maximin score of a also drops by at least (in fact, by exactly) δ. To show that the rest of the proof follows from the outline, we need to show that the maximin score of b drops by at most δ. For each $d \in A \setminus \{b\}$, the weight of the edge from b to d is the sum of fractions of a subset R_d of rankings in π. Now, the collective weight

[10] Recall that the Bucklin score is to be *minimized*.

of rankings in $R_d \cap R_{a \succ c}$ drops by at most δ, and the collective weight of rankings in $R_d \cap R_{c \succ a}$ can only increase. Hence, the weight of each outgoing edge from b drops by at most δ, which means that the maximin score of b also drops by at most δ, as required.

(d) Slater's rule.

Recall that Slater's rule associates a score to every ranking, and then chooses the ranking with the lowest Slater score,[11] breaking ties to choose among all rankings with the lowest Slater score. Even though Slater's rule does not associate scores to alternatives, we show that it fits our framework with a little modification. First, if there are no ties in the unweighted PM graph of a profile π, then similarly to Bucklin's rule, its unweighted PM graph and therefore the Slater scores of all rankings can be preserved in a small enough neighborhood of π, eliminating the possibility of π being a hole. In the general case, we slightly abuse the notation, and use $SC^{\pi}(\sigma)$ to denote the Slater score of ranking σ in profile π.

As in the step 1 of the outline, assume that π is a hole for Slater's rule; the rule returns τ with probability 1 in all interior regions adjacent to x^{π}, but returns a different ranking τ' with a positive probability on π. Then, due to all-inclusivity of the tie-breaking scheme, we must have $SC^{\pi}(\tau') \leqslant SC^{\pi}(\tau)$.[12] We again need to find a σ^* and its associated set of profiles P. P must satisfy all the conditions in the third step of the outline, except we replace the inequality in the scores of alternatives by the inequality in the scores of rankings, namely $SC(\tau') \leqslant SC(\tau)$.

Since $SC^{\pi}(\sigma)$ counts the number of pairwise disagreements of σ with the unweighted PM graph of π, and since small deviations in the fractions x^{π}_{σ} would not change the edges that are not tied, we concentrate on the edges of the PM graph of π that are tied. Formally, let $TIE(\pi)$ denote the set of ordered pairs of alternatives that are tied in the PM graph of π. For $\sigma \in \mathcal{L}(A)$, define

$$s(\sigma) = \sum_{\substack{(c,d) \in TIE(\pi) \\ \text{s.t. } c \succ_{\tau'} d}} \mathbb{I}[c \succ_{\sigma} d] - \sum_{\substack{(c,d) \in TIE(\pi) \\ \text{s.t. } c \succ_{\tau} d}} \mathbb{I}[c \succ_{\sigma} d].$$

It is clear that taking $\sigma^* \in \pi$ such that $s(\sigma) \geqslant 0$ would ensure that in every profile in P, at least as much will be added to the Slater score of τ as to the Slater score of τ' compared to π, ensuring $SC(\tau') \leqslant SC(\tau)$. To see why such a ranking exists, we sum $s(\sigma_i)$ over all votes σ_i in π and, by interchanging the order of summations, we get

$$\begin{aligned}\sum_{i=1}^{n} s(\sigma_i) &= \sum_{\substack{(c,d) \in TIE(\pi) \\ \text{s.t. } c \succ_{\tau'} d}} n^{\pi}_{c \succ d} - \sum_{\substack{(c,d) \in TIE(\pi) \\ \text{s.t. } c \succ_{\tau} d}} n^{\pi}_{c \succ d} \\ &= \frac{n}{2} \cdot \Big(\big|\{(c,d) \in TIE(\pi) \text{ s.t. } c \succ_{\tau'} d\}\big| - \big|\{(c,d) \in TIE(\pi) \text{ s.t. } c \succ_{\tau} d\}\big| \Big)\end{aligned} \tag{7.6}$$

[11] Recall that Slater's score is the disagreement of a ranking from a profile, which must be minimized.

[12] As with Bucklin's rule, the sign of the inequality is reversed because the Slater ranking *minimizes* the Slater score.

$$= 0,$$

where the last step follows since both terms inside the brackets in Equation (7.6) are the number of unordered pairs of alternatives that are tied in the PM graph of π. Hence, there exists a ranking $\sigma^* \in \pi$ with $s(\sigma^*) \geqslant 0$, as required. Finally, $\delta_{\max}$ is chosen so that the non-tied pairs in the PM graph stay non-tied, and the new profile does not fall in an interior region that is not adjacent to x^π.

(e) The ranked pairs method.

This proof does not follow the general outline given above. For an ordered pair of alternatives (c,d), let $w^\pi(c,d)$ denote the weight of the edge from c to d in the weighted PM graph of π. Suppose r outputs a ranking τ with probability 1 in every interior region adjacent to x^π, but does not output τ with probability 1 on π.

Let L denote the list in the ranked pairs process in π where ordered pairs of alternatives are sorted by their weight. Let Δ denote the minimum positive difference between the weights of any two pairs in L. Let (a,b) be the first pair in the list that is chosen with a positive probability and is inconsistent with τ (such a pair exists because r does not output τ with probability 1 on π).

Lemma 7.12. *Let PRE denote the set of pairs in L that have weight strictly greater than the weight of (a,b). Then, each pair in PRE must be chosen with probability 1 or 0 in the ranked pairs process on π, and the subset that is chosen with probability 1 must be consistent with τ.*

Proof. Let L^p be the largest prefix of L such that such that every pair in L^p is chosen with probability 1 or 0 in the ranked pairs process under an inclusive tie-breaking scheme.[13] Let $C^p \subseteq L^p$ be the set of pairs in L^p that are chosen with probability 1.

First, we argue that all pairs in C^p are consistent with τ. Let P denote the space of profiles obtained by changing the fractions of all the rankings by at most $\Delta/(2m!)$. Note that this may only break ties in L, but cannot invert the order of two pairs that were strictly ordered by their weight in L. Similarly to the general outline, P has Hausdorff dimension $k-1$, and hence contains a point in an interior region adjacent to x^π. Further, since ties do not matter for pairs in P, all pairs in P chosen with probability 1 in π would also be chosen with probability 1 in all profiles in P. Hence, all pairs in C^p must be consistent with τ.

Since r does not output τ with probability 1 on π, $L^p \neq L$. Consider the group G of pairs with equal weight that follows L^p. First, G cannot be consistent with C^p, otherwise it would have been part of L^p. Therefore, there must exist a pair $p \in G$ that is chosen with a probability strictly in $(0,1)$ (i.e., not equal to 0 or 1). Thus, there must exist a feasible subset of G such that when it is chosen in the ranked pairs process along with C^p to produce a partial order l, l is inconsistent with p. If l is consistent with τ, then p must be inconsistent with τ. If l is inconsistent with τ, then since C^p is consistent with τ, there must exist a pair in G that is inconsistent with τ.

[13] Note that if L^p has a group of pairs with equal weight, they will all be chosen with probability 1 or all be chosen with probability 0 irrespective of the tie-breaking.

In either case, all pairs in C^p are consistent with τ, and the group G of pairs with equal weight that follows C^p has a pair that is inconsistent with τ. Thus, $(a, b) \in G$, and $PRE = L^p$. ■(Proof of Lemma 7.12)

Next, we argue that $0 < w^\pi(a, b) < 1$. If $w^\pi(a, b) = 0$, then $w^\pi(b, a) = 1$. An ordered pair with weight 1 is consistent with all rankings in the profile. Hence, the set of ordered pairs in π with weight 1 do not contain a cycle. Thus, they are all selected with probability 1 in the ranked pairs process, which is a contradiction as we assumed (a, b) is chosen with a positive probability on π.

On the other hand, if $w^\pi(a, b) = 1$, then all rankings in π must prefer a to b. However, all profiles in $\mathcal{S}(x^\pi)$ have the same set of rankings as π. Hence, the weight of (a, b) is 1 everywhere in $\mathcal{S}(x^\pi)$. Due to the argument presented in the previous paragraph, this implies that in an interior region K adjacent to x^π, a is preferred to b with probability 1. This is a contradiction because r outputs τ with probability 1 in K that prefers b to a.

Hence, indeed $0 < w^\pi(a, b) < 1$. Let $R_{a \succ b}$ be the set of rankings in π that prefer a to b, and let $R_{b \succ a}$ be the set of rankings in π that prefer b to a. Since $0 < w^\pi(a, b) < 1$, we have $R_{a \succ b} \neq \emptyset$ and $R_{b \succ a} \neq \emptyset$.

Recall that Δ is the minimum positive difference between the weights of any two pairs in L. Choose $\delta_{\max} = \Delta/2$. Let P denote the set of profiles obtained by increasing the fractions of rankings in $R_{a \succ b}$ by a total of δ and decreasing the fractions of the rankings in $R_{b \succ a}$ by a total of δ, for $0 < \delta < \delta_{\max}$. This increases the weight of (a, b) by exactly δ and changes the weight of every other pair by at most δ. Due to the choice of $\delta_{\max}$, it is clear that the set of pairs with weight greater than that of (a, b) must be PRE for every profile in P.

Further, the changes in the fractions can only break ties among pairs in PRE, but cannot invert the order of two pairs with different weight in π. Since ties do not matter for pairs in PRE,[14] we see that the same subset of pairs in PRE are chosen in every profile in P. This would imply that under an inclusive tie-breaking scheme, (a, b) has a positive probability of being selected in each profile in P. However, P has Hausdorff dimension $k - 1$, and therefore must contain a point in an interior region K adjacent to x^π. This contradicts the fact that r outputs τ that prefers b to a with probability 1 in K. Hence, π cannot be a hole. ■(Proof of Theorem 7.14)

The comprehensive list of GSRs with no holes includes all prominent rules that are known to be GSRs [146, 204] — suggesting that the no holes property does not impose a significant restriction beyond the assumption that the rule is a GSR. One prominent rule is conspicuously missing — the fascinating but peculiar Dodgson rule [77], which is indeed not a GSR [204].

[14]The pairs in P that were chosen with probability 1 and 0 in π would still be chosen with probability 1 and 0, respectively.

7.8 Impossibility for PM-c and PD-c Rules

Theorem 7.13 establishes the uniqueness of the modal ranking rule within a large family of voting rules (GSRs with no holes). Next we further expand this result by showing that no PM-c or PD-c rule is monotone-robust with respect to all distance metrics. Thus, the modal ranking rule is the unique rule that is monotone-robust with respect to all distance metrics in the union of GSRs with no holes, PM-c rules, and PD-c rules. Crucially, as shown in Figure 7.1, the families of PM-c and PD-c rules are disjoint, and neither one is a strict subset of GSRs.

Theorem 7.15. *For $m \geqslant 3$ alternatives, no PM-c rule or PD-c rule is monotone-robust with respect to all distance metrics.*

Proof. In both parts of this proof (for PM-c rules and PD-c rules), we use an intuitive, but somewhat technical lemma, which is given as Lemma B.2 in the appendix. Let $A = \{a_1, \ldots, a_m\}$ be the set of alternatives. We use $a_{4\ldots m}$ as shorthand for $a_4 \succ \ldots \succ a_m$. Fix $\tau = a_1 \succ \ldots \succ a_m$, and $\sigma^* = a_2 \succ a_1 \succ a_3 \succ a_{4\ldots m}$.

First, we prove that no PM-c rule is monotone-robust with respect to all distance metrics. In particular, using Lemma B.2, we will construct a distance metric d and a d-monotonic noise model G such that no PM-c rule is accurate in the limit for G.

Consider the distribution D over $\mathcal{L}(A)$ defined as follows:

$$
\begin{aligned}
&\Pr_D[a_2 \succ a_1 \succ a_3 \succ a_{4\ldots m}] = \tfrac{4}{9},\\
&\Pr_D[a_1 \succ a_2 \succ a_3 \succ a_{4\ldots m}] = \tfrac{3}{9},\\
&\Pr_D[a_1 \succ a_3 \succ a_2 \succ a_{4\ldots m}] = \tfrac{2}{9},\\
&\Pr_D[\sigma] = 0, \text{ for all } \sigma \text{ not covered above.}
\end{aligned}
$$

By Lemma B.2, we know that there exist a distance metric d and a d-monotonic noise model G such that $\Pr_G[\sigma;\sigma^*] = \Pr_D[\sigma]$ for every $\sigma \in \mathcal{L}(A)$.

Given infinite samples from $G(\sigma^*)$, a $5/9$ fraction — a majority — of the votes have a_1 in the top position. A $7/9$ fraction of the votes prefer a_2 to a_3, while all votes prefer a_2 and a_3 to any other alternative besides a_1. Clearly, a_i is preferred to a_{i+1} for $i \geqslant 4$. Hence, in the PM graph, the alternatives are ordered according to $\tau = a_1 \succ a_2 \succ a_3 \succ a_{4\ldots m}$. Thus, every PM-c rule outputs τ in the limit, which is not the ground truth. Thus, no PM-c rule is accurate in the limit for G.

The construction for PD-c rules is more complex. Here, we will show that there is a noise model such that, given infinite samples for a specific ground truth, the PD graph of the profile induces a ranking that is different from the ground truth. The distribution D above is not sufficient for our purposes since there are pairs of alternatives (e.g., a_2 and a_3) that have the same probability of appearing in the first three positions of the outcome; hence, the PD graph of profiles with infinite samples may not be complete. Instead, we will use a distribution D' so that all probability values of this kind are different.

Let $0 = \delta_1 < \delta_2 < \ldots < \delta_m$ so that $\sum_{i=1}^{m} \delta_i = 1$. Define the probability distribution D'' as follows. Pick one out of the m alternatives so that alternative a_i is picked with probability δ_i. Rank alternative a_i last and complete the ranking by a uniformly random permutation of the alternatives in $\mathcal{L}(A) \setminus \{a_i\}$. Now, the distribution D' is defined as follows: With probability 9/10 (resp., 1/10), the output ranking is sampled from the distribution D (resp., D'').

The important property of distribution D'' is that for every $k \in [m-1]$, the probability that alternative a_i is ranked in the first k positions is exactly $\frac{(1-\delta_i)k}{m-1}$, i.e., strictly decreasing in i. On the other hand, distribution D has the property that for every $k \in [m-1]$, the probability that alternative a_i is ranked in the first k positions is non-increasing in i. Hence, their linear combination D' has the property that for every $k \in [m-1]$, the probability that alternative a_i is ranked in the first k positions is strictly decreasing in i. Additionally,

$$\arg\max_{\tau \in \mathcal{L}(A)} \Pr_{D'}[\tau] = \{\sigma^*\}.$$

Hence, we can apply Lemma B.2 to obtain a distance metric d' and a d'-monotonic noise model G' so that an infinite number of samples from $G'(\sigma^*)$ induce a complete PD graph corresponding to the ranking $\tau = a_1 \succ a_2 \succ a_3 \succ a_{4\ldots m}$, which is different from the ground truth σ^*. Thus, no PD-c rule is accurate in the limit for G'.

We conclude that no PM-c rule or PD-c rule is monotone-robust with respect to all distance metrics. ■ (Proof of Theorem 7.15)

The restriction on the number of alternatives in Theorem 7.15 is indeed necessary. For two alternatives, $\mathcal{L}(A)$ contains only two rankings, and all reasonable voting rules coincide with the majority rule that outputs the more frequent of the two rankings. It can be shown that, in this case, the majority rule is monotone-robust with respect to all distance metrics.

We have shown that the union of PM-c and PD-c rules includes all positional scoring rules, Bucklin's rule, the Kemeny rule, ranked pairs, Copeland's method, and Slater's rule. Two prominent SWFs that are neither PM-c nor PD-c are the maximin rule and STV. In the example given in the proof of Theorem 7.15, the maximin rule and STV would also rank the wrong alternative (a_1) in the first position with probability 1 in the limit. Thus, Theorem 7.15 gives another proof that prominent voting rules are not monotone-robust with respect to all distance metrics.

7.9 Related work

The theme of quantifying the number of samples that are required to uncover the truth plays a central role in a recent paper by Chierichetti and Kleinberg [57]. They study a setting with a single correct alternative and noisy signals about its identity. Focusing on a single voting rule — the plurality rule — they give an upper bound on the number of votes that are required to pinpoint the correct winner. They also prove a lower bound that applies to any voting rule and suggests that plurality is not far from optimal. Interestingly, under Mallows' model we show that plurality is far worse than all PM-c rules, but note that we consider rules that output a ranking while Chierichetti and Kleinberg [57] study rules that output a single winner.

Our initial results regarding the Kemeny rule are related to the work of Braverman and Mossel [41]. Given samples from Mallows' model, they aim to compute the Kemeny ranking; this problem is known to be NP-hard [20]. They focus on circumventing the complexity barrier by giving an efficient algorithm that computes the Kemeny ranking with arbitrarily high probability. In contrast, we ask: How many samples do PM-c rules (including Kemeny) need to reconstruct the *true* ranking?

There is a significant body of literature on MLEs and parameter estimation for noise models over rankings that generalize Mallows' model [71, 92, 129, 137]. In particular, the classic paper by Fligner and Verducci [92] analyzes extensions of Mallows' model with distance functions from two families: those that are based on discordant pairs (including the KT distance) and those that are based on cyclic structure. Critchlow et al. [71] introduce four categories of noise models; they also define desirable axiomatic properties that noise models should satisfy, and determine which properties are satisfied by the different categories. Many papers analyze other random models of preferences, e.g., the Plackett-Luce model [135], the Thurstone-Mosteller model [167], or the random utility model [7].

Somewhat further afield, a recent line of work in computational social choice studies the *distance rationalizability* of voting rules [33, 83, 84, 85, 143]. Voting rules are said to be distance rationalizable if they always select an alternative or a ranking that is "closest" to being a consensus winner, under some notion of distance and some notion of consensus. Among these papers, the one by Elkind et al. [85] is the most closely related to our work; they observe that the Kemeny rule is both an MLE and distance rationalizable via the same distance, and ask whether at least one of several other common rules has the same property (the answer is "no").

Observe that, unlike common voting rules, our modal ranking rule always returns a ranking from the given profile. Endriss and Grandi [87] investigated a similar idea in the context of aggregation of binary opinions where each voter provides a yes/no opinion for a number of issues. In contrast to the modal ranking rule which (in our context) chooses the most frequent vote in the input profile, they investigated rules that return the input vote that is closest to the average or the majority vote.

In our analysis of the modal ranking rule, we technically view the input profile (vector of rankings) as a point in $\mathbb{Q}^{m!}$ ($m!$ is the number of possible rankings), where each coordinate represents the fraction of times a ranking appears in the input profile. This geometric approach to the analysis of voting rules was initiated by Young [208], and was later used by various other authors [70, 146, 159, 187, 188, 205].

7.10 Discussion

While we study three popular distance functions over rankings, we exclude some other distances such as the Cayley distance and the Hamming distance; even the most prominent voting rules such as plurality are not accurate in the limit for any noise model that is monotonic with respect to these distances (see Appendix B.4). On the one hand, this motivates a study of distance functions over rankings that are more appropriate in the social choice context. On the other hand, one may ask: Which voting rules are monotone-robust even with respect to such distance functions?

Furthermore, we have seen that all PM-c rules and all PD-c rules are accurate in the limit for Mallows' model. We later argued that being accurate in the limit for Mallows' model is a very mild requirement, and there are numerous other voting rules that satisfy it. Is it possible to define a much wider class (possibly within the framework of generalized scoring rules [204]) that is accurate in the limit for Mallows' model?

On the conceptual level, we analyze the sample complexity of voting rules as the number of alternatives grows, but our analysis assumes (as is traditionally the case in the literature) that the input to the voting rule is total orders over alternatives. As argued in the introduction, the issue of sample complexity of voting rules directly translates to the problem of estimating the required budget in crowdsourcing tasks. When the number of alternatives is large, obtaining total orders is unrealistic, and inputs with partial information such as pairwise comparisons, partial orders or top-k-lists are employed in practice. Several noise models have been proposed in the literature for the generation of such partial information (see, e.g., [206]). Going one step further, Procaccia et al. [179] proposed a noise model that can incorporate multiple input formats simultaneously given a true underlying ranking. It would be of great practical interest to extend our sample complexity analysis to such noise models.

We mentioned several points of view on the comparison of voting rules: social choice axioms, maximum likelihood estimators, and the distance rationalizability framework. Elkind et al. [85] point out the weakness of the connection between the MLE framework and the DR framework by showing that the Kemeny rule is the only rule that is both MLE and distance rationalizable. We argued that asking for a voting rule to be the *maximum* likelihood estimator is too restrictive, and proposed quantifying the sample complexity instead. This begs the question: How does the relaxed framework of sample complexity relate to the DR framework?

An important conceptual contribution of this chapter is the realization that the modal ranking rule — a natural voting rule that was previously disregarded — is uniquely robust among the union of three large families of voting rules. Figure 7.1 shows a Venn diagram illustrating the relation between these three families. We claim that the modal ranking rule can be exceptionally useful in crowdsourcing settings. Interestingly, from a classic social choice viewpoint the modal ranking rule would appear to be a poor choice. It does satisfy some axiomatic properties, such as Pareto efficiency — if all voters rank x above y, the output ranking places x above y (indeed, the rule always outputs one of the input rankings). But the modal ranking rule fails to satisfy many other basic desiderata, such as monotonicity — if a voter pushes an alternative upwards, and everything else stays the same, that alternative's position in the output should only improve. So our uniqueness result implies an impossibility: a voting rule that is monotone-robust with respect to any distance metric d and is a GSR with no holes, PD-c rule, or PM-c rule, cannot satisfy the monotonicity property. A similar statement is true for any social choice axiom not satisfied by the modal ranking rule. That

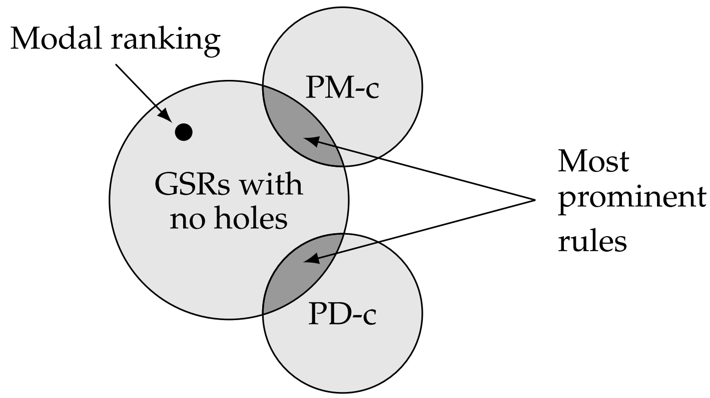

Figure 7.1: The modal ranking rule is uniquely robust within the union of three families of rules.

said, social choice axioms like monotonicity were designed with subjective opinions and notions of social justice in mind. These axioms are incompatible with the settings that motivate our work on a conceptual level, and — as our results show — on a technical level.

Chapter 8

Robust Voting on Social Networks

8.1 Introduction

In the previous chapter, we extended the popular *maximum likelihood estimation* approach to voting by breaking two of its most restrictive requirements. First, we relaxed the assumption that we know the exact noise model that governs the generation of noisy votes given the ground truth, and instead sought *robust voting rules* that provide guarantees with respect to much wider families of noise models. Second, we relaxed the stringent requirement that the voting rule be an MLE by simply requiring the voting rule to be *accurate in the limit* as the number of votes grows.

While this allowed us to obtain promising results, our robust voting framework still carries an unrealistic assumption that is prevalent in most prior work in social choice theory: votes are independent. This assumption is clearly satisfied in some settings, but in many other settings — especially when the voters are people — votes are likely to be correlated through social interactions. We refer to the structure of these interactions as a *social network*, interpreted in the broadest possible sense: any form of interaction qualifies for an edge. From this broad viewpoint, the structure of the social network cannot be known, and, hence, votes are correlated in an unpredictable way. Extending the robustness approach from the previous chapter, we aim to

> *... model the generation of noisy rankings on a social network given a ground truth, and identify voting rules that are accurate in the limit with respect to any network structure and (almost) any choice of model parameters.*

8.2 Our Model

Our starting point is the recently-introduced *independent conversations model* [63]. In this model, there are only two alternatives: one is "correct" (stronger) and one is "incorrect" (weaker). Each edge of the social network is an independent conversation between two voters, whose result (which is independent of the results on other edges — hence the name of the model) is the correct alternative with probability $p > 1/2$, and the incorrect alternative with probability $1 - p$. Then, each voter aggregates the results on the incident edges using the majority rule, and submits the resulting alternative (i.e., the final vote) to the voting rule. Note that if two voters are neighbors in the network, their votes are *not* independent. The voting rule only observes the final votes submitted by the voters (and not the results of conversations on the edges), and must aggregate these votes to find the correct alternative. Conitzer acknowledges that his goal is to "give a simple model that helps to illustrate which phenomena we are likely to encounter as we move to more complex models" [63, p. 1483]. We are indeed interested in a more realistic model that supports multiple alternatives and rests on richer probabilistic foundations.

In our extended model, there is a set of (arbitrarily many) alternatives A. We assume that the voters are connected via an underlying social network structure, represented as an undirected graph $G = (V, E)$ (here, V is the set of voters). We use the notation $e \downarrow v$ to denote that edge e is incident on voter v. Let $E(v) = \{e \in E \mid e \downarrow v\}$ denote the set of edges incident on v, and let $d_v = |E(v)|$ denote the degree of v in the network G. As we explain in Section 8.1, the social network structure may be unknown to us. Our model has four key components.

- *Ground truth.* We assume that each alternative $a \in A$ has a "true quality" denoted by μ_a. The ground truth ranking of the alternatives σ^* ranks the alternatives by their true qualities. We assume that for some constant $\Delta > 0$, we have $|\mu_a - \mu_b| \leqslant \Delta$ for all distinct $a, b \in A$.
- *Quality estimates.* When voters v and v' share an edge, they have an independent discussion. We represent the result of this discussion as a quality estimate for each alternative. Specifically, we associate a random variable $X_{e,a}$ to each edge e for the quality estimate of each alternative a. Crucially, we assume that all $\{X_{e,a}\}_{e \in E, a \in A}$ are mutually independent.
- *Aggregation rules.* We assume that voter v uses an aggregation rule $g_v : \mathbb{R}^{d_v} \to \mathbb{R}$ to derive an aggregate quality estimate $Y_{v,a} = g(\{X_{e,a}\}_{e \in E(v)})$ for each alternative $a \in A$. In the tradition of the random utility theory, his submitted vote σ_v is a ranking of the alternatives by their aggregate quality estimates.
- *Voting rule.* The only information we observe is the set of rankings (votes) submitted by the voters. In particular, we are unaware of the quality estimates sampled on the edges (i.e., values of $X_{e,a}$), or the aggregate quality estimates derived by the voters (i.e., values of $Y_{v,a}$). Moreover, we assume that the distributions of the independent conversations on the edges, the aggregation rules used by the voters, the identities of the voters, and their social network structure are also unknown to us. We use an *anonymous* voting rule $f : \mathcal{L}(A)^n \to \mathcal{L}(A)$ to aggregate the submitted ranked votes into a final ranking of the alternatives. Our goal is to be accurate in the limit, i.e., produce the ground truth ranking σ^* with probability 1 as the number of voters n goes to infinity.

In the next two sections, we instantiate this general model by considering specific distributions of the quality estimates on the edges ($X_{e,a}$) and specific choices of the aggregation rules used by the voters.

8.3 Equal Variance

Let us focus on the following model of independent conversations and aggregation rules.

Quality estimates. Our choice is inspired by the classic *Thurstone-Mosteller model* [148, 194], in which a quality estimate is derived by taking a sample from a Gaussian distribution centered around the true quality. This model is member of the more general class of *random utility models* (see [7] for their use in social choice) in which the distribution need not be Gaussian. In our setting, for each edge $e \in E$ and alternative $a \in A$ we assume $X_{e,a} \sim \mathcal{N}(\mu_a, \nu^2)$, which is a Gaussian distribution with mean μ_a, variance ν^2, and probability density function

$$p(x) = \frac{1}{\sqrt{2\pi\nu^2}} e^{-\frac{(x-\mu_a)^2}{2\nu^2}}.$$

Crucially, we assume that the variance of all the Gaussians is equal, i.e., the noise present in the quality estimates is random noise that is not dependent on the voters or on the alternatives. This is not a weak assumption; we relax it in Section 8.4.

Aggregation rules. We assume that voters aggregate the quality estimates of the alternatives on their incident edges by computing a weighted mean. Specifically, assume that each voter v places a weight $w_v(e) \in \mathbb{R}_{\geqslant 0}$ on each incident edge $e = (v, v') \in E(v)$, which represents how much the voter weights or believes in the conversation with voter v'. Without loss of generality, let the weights be normalized such that $\sum_{e \in E(v)} w_v(e) = 1$ for all $v \in V$. Then, the aggregate quality estimate derived by voter v for alternative a is given by $Y_{v,a} = \sum_{e \in E(v)} w_v(e) X_{e,a}$.

We aim to find voting rules that provide accuracy in the limit for any social network structure G, and for a wide range of choices of the unknown parameters: the true qualities of the alternatives $\{\mu_a\}_{a \in A}$, the variance of the Gaussian distributions ν^2, and the weights assigned by voters to their incident edges $\{w_v(e)\}_{v \in V, e \in E(v)}$. The main difficulty is that the votes of two voters may be correlated when they share an edge in the social network, but the network is unknown to the voting rule. To this end, we first prove a result that shows that under certain conditions, the correlation has negligible effect on the final outcome. We later leverage this result to identify anonymous voting rules that are accurate in the limit.

Lemma 8.1. *Let $Z_v^1, Z_v^2 \in [-\xi, \xi]$ be two bounded random variables associated with each voter $v \in V$, where $\xi > 0$ is a constant. For $i, j \in \{1, 2\}$ and $v, v' \in V$, assume Z_v^i and $Z_{v'}^j$ are independent unless $v = v'$ or $(v, v') \in E$. If there exist positive constants C, γ, δ, and ε such that for all $v \in V$,*

1. *$\mathbb{E}[Z_v^1] - \mathbb{E}[Z_v^2] \geqslant \gamma$, and*
2. *$\Pr[Z_v^1 \leqslant Z_v^2 + \delta] \leqslant C/(d_v)^{1+\varepsilon}$,*

then $\lim_{n \to \infty} \Pr[\sum_{v \in V} Z_v^1 > \sum_{v \in V} Z_v^2] = 1$.

Before we dive into the proof, note that if the random variables were independent, condition 1 and Hoeffding's inequality would have implied the required result. For correlated variables, the intuition is as follows. If d_v is small, then Z_v^1 and Z_v^2 are correlated with only a few other random variables. If d_v is large, then $Z_v^1 > Z_v^2$ holds with high probability anyway. As we later see in Theorem 8.1, this is because voters with large degrees produce accurate votes by assimilating a large amount of independent information from incident edges.

Proof. Partition the set of voters V into two subsets:

$$V_1 = \left\{ v \in V \,\middle|\, d_v \leqslant n^{\frac{1+0.5\cdot\varepsilon}{1+\varepsilon}} \right\} \text{ and } V_2 = V \setminus V_1.$$

Define $Z_{V_i}^j = \sum_{v \in V_i} Z_v^j$ and $Z_V^j = Z_{V_1}^j + Z_{V_2}^j$ for $i, j \in \{1, 2\}$. We wish to prove that $Z_V^1 > Z_V^2$ holds with high probability. We focus on the relations between $Z_{V_1}^1$ and $Z_{V_1}^2$, and between $Z_{V_2}^1$ and $Z_{V_2}^2$ separately, and later combine the two results to prove the required result.

Voters in V_1. Observe that $\mathbb{E}[Z_{V_1}^1 - Z_{V_1}^2] = \sum_{v \in V_1} \mathbb{E}[Z_v^1] - \mathbb{E}[Z_v^2] \geqslant |V_1| \cdot \gamma$. As previously mentioned, we cannot simply use Hoeffding's inequality because the indicator random variables are correlated. We instead use Chebyshev's inequality.

$$\begin{aligned} \Pr[Z_{V_1}^1 \leqslant Z_{V_1}^2] &\leqslant \Pr[|(Z_{V_1}^1 - Z_{V_1}^2) - \mathbb{E}[Z_{V_1}^1 - Z_{V_1}^2]| \geqslant |V_1| \cdot \gamma] \\ &\leqslant \frac{\mathrm{Var}\big(Z_{V_1}^1 - Z_{V_1}^2\big)}{|V_1|^2 \cdot \gamma^2}. \end{aligned} \tag{8.1}$$

Here, $\mathrm{Var}(\cdot)$ denotes the variance of a random variable. To derive an upper bound on $\mathrm{Var}\big(Z_{V_1}^1 - Z_{V_1}^2\big)$, we use the fact that for $i, j \in \{1, 2\}$ and $v, v' \in V_1$, indicator random variables Z_v^i and $Z_{v'}^j$ are only correlated if $v = v'$ or v and v' share an edge (i.e., $(v, v') \in E$). Thus, the random variables corresponding to voter v can be correlated with the random variables corresponding to at most $1 + d_v$ voters. Further, when they are correlated, their covariance satisfies $\mathrm{Cov}(Z_v^i, Z_{v'}^j) \leqslant \sqrt{\mathrm{Var}(Z_v^i) \cdot \mathrm{Var}(Z_{v'}^j)} \leqslant \xi^2$, where the last transition holds because the variance of a $[-\xi, \xi]$-bounded random variable is at most ξ^2 due to Popoviciu's inequality. Hence,

$$\begin{aligned} \mathrm{Var}\big(Z_{V_1}^1 - Z_{V_1}^2\big) &= \sum_{i,j \in \{1,2\}} \sum_{v \in V_1} \sum_{\substack{v' \in V_1: \\ [v'=v] \vee [(v,v') \in E]}} \mathrm{Cov}(Z_v^i, Z_{v'}^j) \\ &\leqslant \xi^2 \cdot \sum_{v \in V_1} 4 \cdot (1 + d_v) \leqslant 4 \cdot \xi^2 \cdot |V_1| \cdot \left(1 + n^{\frac{1+0.5\cdot\varepsilon}{1+\varepsilon}}\right), \end{aligned}$$

where the last transition holds because $d_v \leqslant n^{\frac{1+0.5\cdot\varepsilon}{1+\varepsilon}}$ for all $v \in V_1$. Substituting this into Equation (8.1),

$$\Pr\left[Z_{V_1}^1 \leqslant Z_{V_1}^2\right] \leqslant 4 \cdot \xi^2 \cdot \frac{1 + n^{\frac{1+0.5\cdot\varepsilon}{1+\varepsilon}}}{|V_1| \cdot \gamma^2} \tag{8.2}$$

Note that this probability could be high when $|V_1|$ is small.

Voters in V_2. Fix $v \in V_2$. Then, $d_v \geqslant n^{\frac{1+0.5\cdot\varepsilon}{1+\varepsilon}}$ by the definition of V_2. Hence,

$$\Pr[Z_v^1 \leqslant Z_v^2 + \delta] \leqslant \frac{C}{(d_v)^{1+\varepsilon}} \leqslant \frac{C}{n^{1+0.5\cdot\varepsilon}}, \tag{8.3}$$

where the first transition follows from the second condition assumed in the lemma. Now,

$$\begin{aligned}\Pr[Z_{V_2}^1 \leqslant Z_{V_2}^2 + |V_2| \cdot \delta] &\leqslant \sum_{v \in V_2} \Pr[Z_v^1 \leqslant Z_v^2 + \delta] \\ &\leqslant \frac{C \cdot |V_2|}{n^{1+0.5\cdot\varepsilon}} \leqslant \frac{C}{n^{0.5\cdot\varepsilon}}, \end{aligned} \tag{8.4}$$

where the first transition follows from the Pigeonhole principle, the second transition follows from Equation (8.3), and the last transition holds because $|V_2| \leqslant n$. Note that this probability must go to 0 as $n \to \infty$, unlike $\Pr[Z_{V_1}^1 \leqslant Z_{V_1}^2]$.

We now consider two cases to combine our results.

1. Suppose $|V_2| \geqslant n \cdot 2\xi/(2\xi + \delta)$. Then, $|V_1| \leqslant n \cdot \delta/(2\xi + \delta)$. Observe that we always have $Z_{V_1}^1 - Z_{V_1}^2 \geqslant -|V_1| \cdot 2\xi$. If it holds that $Z_{V_2}^1 - Z_{V_2}^2 > |V_2| \cdot \delta$, then $Z_V^1 > Z_V^2$ follows by adding the two inequalities and substituting the bounds of $|V_1|$ and $|V_2|$. Hence, $\Pr[Z_V^1 \leqslant Z_V^2] \leqslant \Pr[Z_{V_2}^1 \leqslant Z_{V_2}^2 + |V_2| \cdot \delta]$, which goes to 0 as n goes to infinity due to Equation (8.4).
2. Suppose $|V_2| \leqslant n \cdot 2\xi/(2\xi + \delta)$. Then, $|V_1| \geqslant n \cdot \delta/(2\xi + \delta)$. Substituting this into Equation (8.2), we see that $\Pr[Z_{V_1}^1 \leqslant Z_{V_1}^2]$ approaches 0 as n goes to infinity. Equation (8.4) already shows that $\Pr[Z_{V_2}^1 \leqslant Z_{V_2}^2] \leqslant \Pr[Z_{V_2}^1 \leqslant Z_{V_2}^2 + |V_2| \cdot \delta]$ approaches 0 as n goes to infinity. Hence, $\Pr[Z_V^1 \leqslant Z_V^2] \leqslant \Pr[Z_{V_1}^1 \leqslant Z_{V_1}^2] + \Pr[Z_{V_2}^1 \leqslant Z_{V_2}^2]$ goes to 0 as n goes to infinity.

Thus, in both cases we have the desired result. ■

We now use Lemma 8.1 to derive our main result.

Theorem 8.1. *If there exists a universal constant $D \in \mathbb{N}$ such that $\sum_{e \in E(v)} [w_v(e)]^2 \leqslant \Delta^2/(8\,\nu^2 \ln d_v)$ for all voters v with degree $d_v \geqslant D$, then all PM-c rules, the modal ranking rule, and all strict positional scoring rules are accurate in the limit irrespective of the choices of the unknown parameters: the social network structure G, the true qualities $\{\mu_a\}_{a \in A}$, the variance ν^2, and the weights $\{w_v(e)\}_{v \in V, e \in E(v)}$.*

Before we prove the result, we remark that the bound on $\sum_{e \in E(v)} [w_v(e)]^2$ is a mild restriction. In our setting with normalized weights $\left(\sum_{e \in E(v)} w_v(e) = 1\right)$, the unweighted mean has $\sum_{e \in E(v)} [w_v(e)]^2 = 1/d_v$ which is much smaller than our required bound. More generally, the condition is satisfied if no voter v places an excessive weight — specifically, a weight greater than $\Delta/(4\nu\sqrt{d_v \ln d_v})$ — on any single incident edge.

Proof of Theorem 8.1. Let us begin with PM-c rules.

PM-c Rules. Recall that PM-c rules are guaranteed to return the ground truth ranking σ^* if the pairwise majority graph is consistent with σ^*. We wish to use Lemma 8.1 to show that for every pair of alternatives $a, b \in A$ such that $a \succ_{\sigma^*} b$, there would be an edge from a to b in the pairwise majority graph of the profile consisting of the votes submitted by the voters with probability 1 as n goes to infinity. Applying the union bound over all pairs of alternatives implies that the entire pairwise majority graph would be consistent with σ^* with probability 1 as n goes to infinity.

Now, for voter $v \in V$ and alternative $a \in A$, the aggregate quality estimate $Y_{v,a} = \sum_{e \in E(v)} w_v(e) X_{e,a}$ follows the distribution $\mathcal{N}(\mu_a, \nu^2 \sum_{e\in E(v)} [w_v(e)]^2)$ because each $X_{e,a} \sim \mathcal{N}(\mu_a, \nu^2)$. Let $(W_v)^2 = \sum_{e\in E(v)} [w_v(e)]^2$.

Fix alternatives $a, b \in A$ such that $a \succ_{\sigma^*} b$ (thus, $\mu_a > \mu_b$). Note that $Y_{v,a} - Y_{v,b} \sim \mathcal{N}(\mu_a - \mu_b, 2\,\nu^2\,(W_v)^2)$. Now, recall that there is an edge from a to b in the pairwise majority graph if a strict majority of the voters prefer a to b, i.e., if $\sum_{v \in V} \mathbb{I}[Y_{v,a} > Y_{v,b}] > n/2$ (where $\mathbb{I}$ is the indicator random variable). Hence, in Lemma 8.1 we take $Z_v^1 = \mathbb{I}[Y_{v,a} > Y_{v,b}]$ and $Z_v^2 = \mathbb{I}[Y_{v,a} \leqslant Y_{v,b}]$. Finally, we complete the proof by showing that the two conditions required by Lemma 8.1 hold.

Condition 1*:* $\mathbb{E}[Z_v^1] - \mathbb{E}[Z_v^2] \geqslant \gamma$, where $\gamma > 0$ is a constant. Note that $\mathbb{E}[Z_v^1] - \mathbb{E}[Z_v^2] = 2 \cdot \Pr[Y_{v,a} > Y_{v,b}] - 1$. Since $Y_{v,a} - Y_{v,b} \sim \mathcal{N}(\mu_a - \mu_b, 2\,\nu^2\,(W_v)^2)$, we have that

$$\Pr[Y_{v,a} - Y_{v,b} > 0] = \Phi\left(\frac{\mu_a - \mu_b}{\sqrt{2}\,\nu\, W_v}\right) \geqslant \Phi\left(\frac{\Delta}{\sqrt{2}\,\nu}\right) \geqslant \frac{1}{2} + \gamma',$$

where $\gamma' > 0$ is a constant. Here, the second transition holds because $W_v \leqslant 1$ and $\mu_a - \mu_b \geqslant \Delta$, and the final transition is a standard property of the Gaussian distribution. Hence, condition 1 holds with $\gamma = 2\gamma'$.

Condition 2*:* $\Pr[Z_v^1 \leqslant Z_v^2 + \delta] = O(1/(d_v)^{1+\varepsilon})$, where $\varepsilon, \delta > 0$ are constants. Take $\delta = 0.5$, and recall that Z_v^1 and Z_v^2 are indicator random variables. Then,

$$\Pr[Z_v^1 \leqslant Z_v^2 + \delta] = \Pr[Z_v^1 = 0 \vee Z_v^2 = 1] = \Pr[Y_{v,a} \leqslant Y_{v,b}].$$

Since $Y_{v,a} - Y_{v,b} \sim \mathcal{N}(\mu_a - \mu_b, 2\,\nu^2\,(W_v)^2)$, we have

$$\Pr[Y_{v,a} - Y_{v,b} \leqslant 0] = 1 - \Phi(\lambda) \leqslant \frac{1}{\sqrt{2\pi} \cdot \lambda} e^{-\lambda^2/2},$$

where $\lambda = (\mu_a - \mu_b)/(\sqrt{2}\,\nu\, W_v)$, and the last transition is a standard upper bound for Gaussian distributions. Substituting our assumption that $(W_v)^2 \leqslant \Delta^2/(8\,\nu^2\, \ln d_v)$ and simplifying, we obtain that the probability is $O(1/(d_v)^2)$. Hence, condition 2 holds with $\varepsilon = 1$.

Since both conditions are satisfied, Lemma 8.1 implies that every PM-c rule is accurate in the limit.

Modal Ranking Rule. Recall that the modal ranking rule chooses the most frequent ranking in the profile. Thus, we need to show that the ground truth ranking appears more frequently than any other ranking. Fix a ranking $\sigma \neq \sigma^*$, and define $Z_v^1 = \mathbb{I}[\sigma_v = \sigma^*]$ and $Z_v^2 = \mathbb{I}[\sigma_v = \sigma]$. Then, we wish to use Lemma 8.1 to show that the number of occurrences of σ^* in the profile is larger than the number of occurrences of σ with probability 1 as n goes to infinity. Applying the union bound over all rankings $\sigma \neq \sigma^*$ would imply that σ^* would be the most frequent ranking in the profile with probability 1 as n goes to infinity. Thus, the modal ranking rule would be accurate in the limit.

Next, we show that the two conditions of Lemma 8.1 hold.

Condition 1: $\mathbb{E}[Z_v^1] - \mathbb{E}[Z_v^2] \geqslant \gamma$, where $\gamma > 0$ is a constant. To derive this, we leverage a result by Jiang et al. [115]. Using techniques from the proof of their Theorem 2, it can be shown that if we obtain a ranking σ by sampling utilities from Gaussians and ordering the alternatives by their sampled utilities, then for any ranking $\tau \in \mathcal{L}(A)$ and alternatives $a, b \in A$ such that $a \succ_{\sigma^*} b$ and $a \succ_\tau b$, we have $\Pr[\sigma = \tau] - \Pr[\sigma = \tau_{a\leftrightarrow b}]$ is at least a positive constant, where $\tau_{a\leftrightarrow b}$ denotes the ranking obtained by swapping alternatives a and b in τ. That is, swapping two alternatives to match their order as in σ^* increases the probability of the ranking being sampled by at least a positive constant. However, this result uses a lower bound on the variances of the Gaussian distributions from which quality estimates are sampled. In our case, no such lower bound may exist for vertices with high degree. However, in the absence of such a lower bound one can still show that for the ranking σ_v of voter v, we have that $\Pr[\sigma_v = \tau] - \Pr[\sigma_v = \tau_{a\leftrightarrow b}]$ is non-negative for every $\tau \in \mathcal{L}(A)$ (with $a \succ_\tau b$), and is at least a positive constant γ' when $\tau = \sigma^*$. This is presented as Lemma C.1 in Appendix C.1.

Finally, to show that $\Pr[\sigma_v = \sigma^*] - \Pr[\sigma_v = \sigma] \geqslant \gamma$ (where $\gamma > 0$ is a constant), we start from ranking σ and perform "bubble sort" to convert it into σ^*. That is, in each iteration we find a pair that is ordered differently than in σ^*, and swap the pair. Note that this process converges to σ^* in at most m^2 iterations, and the probability of the ranking never decreases, and increases by at least γ' in the last iteration. This proves that condition 1 holds with $\gamma = \gamma'$.

Condition 2: $\Pr[Z_v^1 \leqslant Z_v^2 + \delta] = O(1/(d_v)^{1+\varepsilon})$ for constants $\delta, \varepsilon > 0$. This condition is very easy to establish. Again, take $\delta = 0.5$. Then,

$$\begin{aligned}\Pr[Z_v^1 \leqslant Z_v^2 + \delta] &= \Pr[Z_v^1 = 0 \vee Z_v^2 = 1] = \Pr[\sigma_v \neq \sigma^*] \\ &\leqslant \sum_{a,b \in A \,:\, a \succ_{\sigma^*} b} \Pr[Y_{v,a} \leqslant Y_{v,b}],\end{aligned}$$

where the last transition holds because if σ_v does not match σ^*, then there exist alternatives a and b such that $a \succ_{\sigma^*} b$ but $b \succ_{\sigma_v} a$ (thus, $Y_{v,a} \leqslant Y_{v,b}$). However, note that this probability is at most m^2 times the probability obtained in condition 2 for PM-c rules, which was $O(1/(d_v)^2)$. Because the number of alternatives m is a constant in our model, multiplying by m^2 does not increase the order in terms of d_v. Hence, condition 2 also holds with $\varepsilon = 1$.

In conclusion, Lemma 8.1 implies that the modal ranking rule is accurate in the limit, as required.

PD-c Rules. We want to show that strict scoring rules are accurate in the limit. Take a strict scoring rule with score vector $(\alpha_1, \ldots, \alpha_m)$. Recall that $\alpha_i > \alpha_{i+1}$ for all $i \in$

$\{1,\ldots,m-1\}$. We use Lemma 8.1 to show that for every pair of alternatives $a,b \in A$ with $a \succ_{\sigma^*} b$, the score of a is greater than the score of b with probability 1 in the limit as the number of voters goes to infinity. Then, applying the union bound over all pairs of alternatives would yield the desired result.

Fix $a,b \in A$ with $a \succ_{\sigma^*} b$. Let Z_v^1 and Z_v^2 denote the scores given by voter v to alternatives a and b, respectively. Note that Z_v^1 and Z_v^2 are bounded random variables. Then, Z_V^1 and Z_V^2 denote the overall scores of a and b, respectively. Next, we show that the two conditions of Lemma 8.1 hold.

Condition 1: $\mathbb{E}[Z_v^1] - \mathbb{E}[Z_v^2] \geqslant \gamma$, where $\gamma > 0$ is a constant. Note that

$$\begin{aligned} &\mathbb{E}[Z_v^1] - \mathbb{E}[Z_v^2] \\ &= \sum_{i=1}^{m} \alpha_i \cdot (\Pr[\sigma_v(a) = i] - \Pr[\sigma_v(b) = i]) \\ &= \sum_{j=1}^{m-1} (\alpha_j - \alpha_{j+1}) \cdot \sum_{i=1}^{j} (\Pr[\sigma_v(a) = i] - \Pr[\sigma_v(b) = i]) \end{aligned} \tag{8.5}$$

Let us denote $\{1,\ldots,j\}$ by $[j]$. Then, $\mathbb{E}[Z_v^1] - \mathbb{E}[Z_v^2] = \sum_{j=1}^{m-1}(\alpha_j - \alpha_{j+1}) \cdot (\Pr[\sigma_v(a) \in [j]] - \Pr[\sigma_v(b) \in [j]])$.

Jiang et al. [115] showed that $\Pr[\sigma_v(a) \in [j]] - \Pr[\sigma_v(b) \in [j]]$ is at least a constant for every $j \in \{1,\ldots,m-1\}$. However, they assume a constant *lower* and upper bound on the variance. When the variance can be arbitrarily low, $\sigma_v(a)$ and $\sigma_v(b)$ would coincide with $\sigma^*(a)$ and $\sigma^*(b)$, respectively, with very high probability. Thus, we cannot expect $\Pr[\sigma_v(a) \in [j]] - \Pr[\sigma_v(b) \in [j]]$ to be at least a constant for every $j \in \{1,\ldots,m-1\}$. Instead, we can show that in our case with only a constant upper bound on the variance, the difference is non-negative for every $j \in \{1,\ldots,m-1\}$, and at least a constant for some $j \in \{1,\ldots,m-1\}$. This is presented as Lemma C.2 in Appendix C.1. Note that this is sufficient to show that in Equation (8.5), $\mathbb{E}[Z_v^1] - \mathbb{E}[Z_v^2]$ is at least a positive constant because $\alpha_j - \alpha_{j+1}$ is at least a positive constant for every $j \in \{1,\ldots,m-1\}$.

Condition 2: $\Pr[Z_v^1 \leqslant Z_v^2 + \delta] = O(1/(d_v)^{1+\varepsilon})$ for constants $\delta, \varepsilon > 0$. While establishing condition 2 for the modal ranking rule, we proved that $\Pr[\sigma_v \neq \sigma^*] = O(1/(d_v)^2)$. When $\sigma_v = \sigma^*$, then $Z_v^1 = \alpha_{\sigma^*(a)}$ and $Z_v^2 = \alpha_{\sigma^*(b)}$. Hence, $Z_v^1 > Z_v^2 + \delta$ for constant $\delta = (1/2) \cdot \min_{j \in \{1,\ldots,m-1\}} \alpha_j - \alpha_{j+1}$. Thus, $\Pr[Z_v^1 \leqslant Z_v^2 + \delta] \leqslant \Pr[\sigma_v \neq \sigma^*] = O(1/(d_v)^{1+\varepsilon})$ is satisfied with $\varepsilon = 1$. ■

While all strict positional scoring rules are accurate in the limit irrespective of the social network structure, one can show that other positional scoring rules such as plurality are not always accurate in the limit; an example is presented in Appendix C.2.

8.4 Unequal Variance

In the previous section we showed that PM-c rules, the modal ranking rule, and strict scoring rules are accurate in the limit when the independent conversations on the edges produce quality estimates from Gaussian distributions (with equal variance) and voters aggregate them using a weighted mean. The equal variance assumption is perhaps the most restrictive assumption in the model of Section 8.3. In this section, we analyze a more general model, which is identical to the model of Section 8.3, except for allowing Gaussians with different variance. Formally, we instantiate our general model using the following model of quality estimates.

Quality estimates. For each edge $e \in E$ and alternative $a \in A$, assume $X_{e,a} \sim \mathcal{N}(\mu_a, (\nu_{e,a})^2)$. Crucially, we assume that all $(\nu_{e,a})^2$ are upper bounded by a global constant. For notational convenience, denote this constant by ν^2. Hence, $(\nu_{e,a})^2 \leqslant \nu^2$ for all $e \in E$ and $a \in A$.

Computer-based simulations provided non-trivial counterexamples (presented in Appendix C.3) showing that unequal variance invalidates Theorem 8.1 with respect to strict positional scoring rules and the modal ranking rule.

Theorem 8.2. *There exist a social network graph $G = (V, E)$, true qualities of alternatives $\{\mu_a\}_{a \in A}$, and Gaussian random variables $X_{e,a}$ for each edge $e \in E$ and alternative $a \in A$ whose variances depend on the alternative a, for which the modal ranking rule is not accurate in the limit, and there exists a strict scoring rule (in particular, Borda count) which is not accurate in the limit.*

In a nutshell, the key insight is that we find a Gaussian distribution for each alternative $a \in A$ such that ranking the alternatives based on a quality estimate sampled from their Gaussian distribution leads to: (i) a ranking other than the true ranking is returned with a probability higher than that of the true ranking itself, which causes the modal ranking rule to fail to achieve accuracy in the limit, and (ii) the probabilities of different alternatives being placed in various positions is such that between two alternatives, the less preferred alternative in the true ranking has greater expected Borda score than the more preferred alternative, causing Borda count to violate accuracy in the limit. Despite these counterintuitive phenomena, it holds that the top alternative in the true ranking is ranked higher than the alternative ranked second in the true ranking with probability strictly greater than $1/2$, and a similar statement also holds for all other pairs of alternatives, thereby ensuring that PM-c rules are accurate in the limit.

Happily, the success of PM-c rules is not a coincidence. Indeed, note that in our proof of Theorem 8.1 we leverage the results of Jiang et al. [115] to prove condition 1 of Lemma 8.1 for PD-c rules and the modal ranking rule. Jiang et al. crucially assume that all distributions have equal variance, and their results break down when this assumption is violated. On the other hand, our proof for the accuracy in the limit of PM-c rules does not rely on their results, and, in fact, does not make use of the equal variance assumption. Specifically, with unequal variance we have that $Y_{v,a} - Y_{v,b}$ follows the Gaussian distribution $\mathcal{N}(\mu_a - \mu_b, \sum_{e \in E(v)} [w_v(e)]^2 \cdot ((\nu_{e,a})^2 + (\nu_{e,b})^2))$. Note that our proof only uses an upper bound on the variance of this Gaussian distribution, and the variance is still upper bounded by $2\,\nu^2\,(W_v)^2$. Hence, for PM-c rules, the proof of Theorem 8.1 goes through even with unequal variance, and shows that all PM-c rules are accurate in the limit.

Theorem 8.3. *Assume that there exists a constant ν such that $(\nu_{e,a})^2 \leqslant \nu^2$ for all $e \in E$ and $a \in A$, and a universal constant $D \in \mathbb{N}$ such that $\sum_{e\in E(v)}[w_v(e)]^2 \leqslant \Delta^2/(8\,\nu^2 \ln d_v)$ for all voters v with degree $d_v \geqslant D$. Then, all PM-c rules are accurate in the limit irrespective of the choices of the unknown parameters: the true qualities $\{\mu_a\}_{a\in A}$, the variances $\{(\nu_{e,a})^2\}_{e\in E,a\in A}$, and the weights $\{w_v(e)\}_{v\in V,e\in E(v)}$.*

Theorem 8.3 establishes that *PM-c rules are qualitatively more robust than PD-c rules and the modal ranking rule* in our setting: While PD-c rules and the modal ranking rule lose their accuracy in the limit when relaxing the equal variance assumption, PM-c rules still guarantee accuracy in the limit irrespective of all unknown parameters. In fact, observe that in our proof for PM-c rules, we only require that for every pair of alternatives $a, b \in A$ with $a \succ_{\sigma^*} b$, we have (i) $\Pr[Y_{v,a} > Y_{v,b}] > 1/2$, and (ii) both $Y_{v,a}$ and $Y_{v,b}$ are *sufficiently concentrated* around their respective means μ_a and μ_b so that $\Pr[Y_{v,a} \leqslant Y_{v,b}] = o(1/d_v)$. Using this observation, we can extend the robustness of PM-c rules beyond the restrictions imposed by Theorem 8.3 in both dimensions: the possible distributions on the edges and the possible aggregation rules used by the voters.

For example, leveraging an elegant extension of the classic McDiarmid inequality by Kontorovich [124], we can show that PM-c rules are accurate in the limit when the distributions on the edges have finite "subgaussian diameter" (this includes all distributions with bounded support and all Gaussian distributions) and voters use weighted mean aggregation. On the other hand, using a concentration inequality for medians, one can show that when the distributions on the edges are Gaussians with bounded variance, then the voters could also use weighted median (instead of weighted mean) aggregation, and PM-c rules would remain accurate in the limit.

8.5 Related Work

The work presented in this chapter is closely related to two papers by Conitzer [62, 63]. The independent conversations model of the latter paper was discussed above. Importantly, the challenge Conitzer [63] addresses is quite different from ours: he is interested in finding the *maximum likelihood estimator (MLE)* for the ground truth, i.e., he wants to know which of the two alternatives is more likely to be correct, given the observed (binary) votes. The answer strongly depends on social network structure, and his main result is that, in fact, the problem is #$\mathcal{P}$-hard. In an earlier, brief note, Conitzer [62] is also interested in the maximum likelihood approach to noisy voting on a social network. While the model he introduces also extends to the case of more than two alternatives, the assumptions of the model are such that the (known) network structure is essentially irrelevant, that is, the maximum likelihood estimator is invariant to network structure.

While the above papers are, to our knowledge, the only papers that deal with the MLE approach to voting on a social network, there is a substantial body of work on the MLE approach to voting more generally [7, 8, 9, 65, 70, 86, 136, 179, 206, 207, 209]. However, all of these papers assume that votes are drawn i.i.d. (conditional on the true ranking) from a noise model.

A bit further afield, there is a large body of work that studies the diffusion of opinions, votes, technologies, or products (but not ranked estimates) in a social network. An especially pertinent example is the work of Mossel et al. [147], where at each time step voters adopt the most popular opinion among their neighbors, and at some point opinions are aggregated via the plurality rule. Other popular diffusion models include the *independent cascade model*, the *linear threshold model*, and the DeGroot model [72]; see the survey by Kleinberg [122] for a fascinating overview.

8.6 Discussion

Let us briefly discuss several pertinent issues.

Temporal dimension. While in our model each voter performs a one-time, synchronous aggregation of information from its incident edges, in general voters may perform multiple and/or asynchronous updates. After k updates, the information possessed by a voter would be a weighted aggregation of the information from all nodes up to distance k from the voter, although the weight associated with another voter at distance k would presumably be exponentially small in k. Deriving positive robustness results in this model seems to require making our simple covariance bounds more sensitive to weights. We believe that Gaussian hypercontractivity results [145] may be helpful in this context.

Opinions on vertices. The *independence* part of our extension of the independent conversations model seems to be a restrictive assumption because the conversations of a voter with two other voters are likely to be positively correlated through the (prior) beliefs of the voter. In this sense, it seems more natural to consider a model where the opinions are attached to vertices rather than edges. Specifically, one might consider a model where the prior opinion of each voter is first drawn from a distribution, and then voters are allowed to aggregate opinions from their neighbors. This leads to immediate impossibilities. Indeed, consider a star network where all peripheral voters give weight 1 to the central voter and 0 to themselves (this does not violate the conditions of Theorem 8.1). At the end, all peripheral voters would have perfectly correlated votes, coinciding with the prior opinion of the central voter which is inaccurate with a significant probability. It follows that any reasonable anonymous voting rule, which would output this opinion, would not be accurate in the limit. Interestingly, we can circumvent this impossibility easily if we know the social network structure: We can simply return the vote submitted by the central voter, which is guaranteed to be accurate as the central voter assimilates information from many sources.

Ground truth and opinion formats. Finally, we assume that the ground truth is a true quality for each alternative, which leads us to a random utility based model. Another compelling alternative is to assume that the ground truth is only an ordinal ranking of the alternatives. In this case, the samples on the edges would also be rankings (instead of noisy quality estimates), and voters would aggregate rankings on their incident edges using their own local voting rules. This model gives rise to many counterintuitive phenomena. For example, using Borda count to aggregate two rankings sampled from the popular Mallows' model [139] with noise parameters $\varphi = 0.1$ and $\varphi = 0.9$ leads to a ranking that is not the ground truth being returned with higher probability than the ground truth itself, ultimately showing that Borda count would not be accurate in the limit. Remarkably, popular PM-c rules seem to be robust against such examples, hinting at the possibility that PM-c rules may also possess compelling robustness properties in this model.

Chapter 9

A Worst-Case Approach to Voting

9.1 Introduction and Our Approach

In Chapters 7 and 8, we developed the robustness approach to voting, which strengthens the traditional MLE approach by designing voting rules that provide guarantees with respect to not just a single noise model, but a wide family of noise models. This required us to ask for weaker guarantees such as accuracy in the limit. While the rules designed are relevant for large-scale applications, they may not be appropriate when the number of votes is not large. In fact, even when the number of votes is extremely small, one should be able to find a good estimate of the ground truth given that each vote is highly accurate. However, the statistical framework of voting (including both the robustness approach and the MLE approach) fail to provide guarantees in this case.

In this chapter, we propose a fundamentally different approach to aggregating noisy votes, which alleviates all of the aforementioned concerns. Instead of assuming probabilistic noise, we assume a known upper bound on the "total noise" in the input votes, and allow the input votes to be *adversarial* subject to the upper bound. We emphasize that in potential application domains there is no adversary that actively inserts errors into the votes; we choose an adversarial error model to be able to correct errors even in the *worst case*. This style of worst-case analysis — where the worst case is assumed to be generated by an adversary — is prevalent in many branches of computer science, e.g., in the analysis of online algorithms [32], and in machine learning [26, 117].

We wish to design voting rules that do well in this worst-case scenario. From this viewpoint, our approach is closely related to the extensive literature on *error-correcting codes*. One can think of the votes as a *repetition code*: each vote is a transmitted noisy version of a "message" (the ground truth). The task of the "decoder" is to correct adversarial noise and recover the ground truth, given an upper bound on the total error. The question is: how much total error can this "code" allow while still being able to recover the ground truth?

In more detail, let d be a distance metric on the space of rankings. As an example, recall that the Kendall tau (KT) distance between two rankings measures the number of pairs of alternatives on which the two rankings disagree. Suppose that we receive

n votes over the set of alternatives $\{a, b, c, d\}$, for an even n, and we know that the average KT distance between the votes and the ground truth is at most $1/2$. Can we always recover the ground truth? No: in the worst-case, exactly $n/2$ agents swap the two highest-ranked alternatives and the rest report the ground truth. In this case, we observe two distinct rankings (each $n/2$ times) that only disagree on the order of the top two alternatives. Both rankings have an average distance of $1/2$ from the input votes, making it impossible to determine which of them is the ground truth.

Let us, therefore, cast a larger net. Inspired by *list decoding* of error-correcting codes (see, e.g., [107]), our main research question is:

> *Fix a distance metric d. Suppose that we are given n noisy rankings, and that the average distance between these rankings and the ground truth is at most t. We wish to recover a ranking that is guaranteed to be at distance at most k from the ground truth. How small can k be, as a function of n and t?*

9.2 Preliminaries

As in the previous chapters, let A be the set of alternatives ($m = |A|$), and $\mathcal{L}(A)$ be the set of rankings over A. A profile $\pi \in \mathcal{L}(A)^n$ is a collection of n votes (rankings). We are interested in voting rules (technically, social welfare functions) $f : \mathcal{L}(A)^n \to \mathcal{L}(A)$ that map every profile to a ranking. Once again, following the notations from the previous chapters, we assume that there exists an underlying ground truth ranking $\sigma^* \in \mathcal{L}(A)$ of the alternatives.

We use a distance metric d over $\mathcal{L}(A)$ to measure errors; the error of a vote σ with respect to σ^* is $d(\sigma, \sigma^*)$, and the average error of a profile π with respect to σ^* is $d(\pi, \sigma^*) = (1/n) \cdot \sum_{\sigma \in \pi} d(\sigma, \sigma^*)$. We consider four popular distance metrics over rankings in this chapter: the Kendall tau (KT) distance d_{KT}, the (Spearman's) footrule (FR) distance d_{FR}, the maximum displacement (MD) distance d_{MD}, and the Cayley (CY) distance d_{CY}. The first three distance metrics are defined in Chapter 7. We define the Cayley distance below.

- *The Cayley (CY) distance*, denoted d_{CY}, measures the minimum number of swaps (not necessarily of adjacent alternatives) required to convert one ranking into another.

All four metrics described above are *neutral*: A distance metric is called neutral if the distance between two rankings is independent of the labels of the alternatives; in other words, choosing a relabeling of the alternatives and applying it to two rankings keeps the distance between them invariant.

9.3 Worst-Case Optimal Rules

Suppose we are given a profile π of n noisy rankings that are estimates of an underlying true ranking σ^*. In the absence of any additional information, any ranking could potentially be the true ranking. However, because essentially all crowdsourcing methods draw their power from the often-observed fact that individual opinions are accurate on average, we can plausibly assume that while some agents may make many mistakes, the average error is fairly small. An upper bound on the average error may be inferred by observing the collected votes, or from historical data (but see the next section for the case where this bound is inaccurate).

Formally, suppose we are guaranteed that the average distance between the votes in π and the ground truth σ^* is *at most* t according to a metric d, i.e., $d(\pi, \sigma^*) \leqslant t$. With this guarantee, the set of possible ground truths is given by the "ball" of radius t around π.

$$\mathcal{B}_t^d(\pi) = \{\sigma \in \mathcal{L}(A) \mid d(\pi, \sigma) \leqslant t\}.$$

Note that we have $\sigma^* \in \mathcal{B}_t^d(\pi)$ given our assumption; hence, $\mathcal{B}_t^d(\pi) \neq \emptyset$. We wish to find a ranking that is as close to the ground truth as possible. Since our approach is worst case in nature, our goal is to find the ranking that minimizes the maximum distance from the possible ground truths in $\mathcal{B}_t^d(\pi)$. For a set of rankings $S \subseteq \mathcal{L}(A)$, let its *minimax ranking*, denoted $\text{MINIMAX}^d(S)$, be defined as follows.[1]

$$\text{MINIMAX}^d(S) = \underset{\sigma \in \mathcal{L}(A)}{\arg\min} \max_{\sigma' \in S} d(\sigma, \sigma').$$

Let the *minimax distance* of S, denoted $k^d(S)$, be the maximum distance of $\text{MINIMAX}^d(S)$ from the rankings in S according to d. Thus, given a profile π and the guarantee that $d(\pi, \sigma^*) \leqslant t$, the worst-case optimal voting rule OPT^d returns the minimax ranking of the set of possible ground truths $\mathcal{B}_t^d(\pi)$. That is, for all profiles $\pi \in \mathcal{L}(A)^n$ and $t > 0$,

$$\text{OPT}^d(t, \pi) = \text{MINIMAX}^d\left(\mathcal{B}_t^d(\pi)\right).$$

Furthermore, the output ranking is guaranteed to be at distance at most $k^d(\mathcal{B}_t^d(\pi))$ from the ground truth. We overload notation, and denote $k^d(t, \pi) = k^d(\mathcal{B}_t^d(\pi))$, and

$$k^d(t) = \max_{\pi \in \mathcal{L}(A)^n} k^d(t, \pi).$$

While k^d is explicitly a function of t, it is also implicitly a function of n. Hereinafter, we omit the superscript d whenever the metric is clear from context. Let us illustrate our terminology with a simple example.

Example 9.1. Let $A = \{a, b, c\}$. We are given profile π consisting of 5 votes: $\pi = \{2 \times (a \succ b \succ c), a \succ c \succ b, b \succ a \succ c, c \succ a \succ b\}$.

The maximum distances between rankings in $\mathcal{L}(A)$ allowed by d_{KT}, d_{FR}, d_{MD}, and d_{CY} are 3, 4, 2, and 2, respectively; let us assume that the average error limit is half the maximum distance for all four metrics.[2]

Consider the Kendall tau distance with $t = 1.5$. The average distances of all 6 rankings from π are given below.

[1]We use $\text{MINIMAX}^d(S)$ to denote a single ranking. Ties among multiple minimizers can be broken arbitrarily; our results are independent of the tie-breaking scheme.

[2]Scaling by the maximum distance is not a good way of comparing distance metrics; we do so for the sake of illustration only.

$d_{KT}(\pi, a \succ b \succ c) = 0.8$	$d_{KT}(\pi, a \succ c \succ b) = 1.0$
$d_{KT}(\pi, b \succ a \succ c) = 1.4$	$d_{KT}(\pi, b \succ c \succ a) = 2.0$
$d_{KT}(\pi, c \succ a \succ b) = 1.6$	$d_{KT}(\pi, c \succ b \succ a) = 2.2$

Thus, the set of possible ground truths is $\mathcal{B}_{1.5}^{d_{KT}}(\pi) = \{a \succ b \succ c,\ a \succ c \succ b,\ b \succ a \succ c\}$. This set has a unique minimax ranking $\text{OPT}^{d_{KT}}(1.5, \pi) = a \succ b \succ c$, which gives $k^{d_{KT}}(1.5, \pi) = 1$. Table 9.1 lists the sets of possible ground truths and their minimax rankings[3] under different distance metrics.

Voting Rule	Possible Ground Truths $\mathcal{B}_t^d(\pi)$	Output Ranking
$\text{OPT}^{d_{KT}}(1.5, \pi)$, $\text{OPT}^{d_{CY}}(1, \pi)$	$\left\{\begin{array}{c} a \succ b \succ c, \\ a \succ c \succ b, \\ b \succ a \succ c \end{array}\right\}$	$a \succ b \succ c$
$\text{OPT}^{d_{FR}}(2, \pi)$, $\text{OPT}^{d_{MD}}(1, \pi)$	$\left\{\begin{array}{c} a \succ b \succ c, \\ a \succ c \succ b \end{array}\right\}$	$\left\{\begin{array}{c} a \succ b \succ c, \\ a \succ c \succ b \end{array}\right\}$

Table 9.1: Application of the optimal voting rules on π.

Note that even with identical (scaled) error bounds, different distance metrics lead to different sets of possible ground truths as well as different optimal rankings. This demonstrates that the choice of the distance metric is important.

9.3.1 Upper Bound

Given a distance metric d, a profile π, and that $d(\pi, \sigma^*) \leqslant t$, we can bound $k(t, \pi)$ using the *diameter* of the set of possible ground truths $\mathcal{B}_t(\pi)$. For a set of rankings $S \subseteq \mathcal{L}(A)$, denote its diameter by $\mathcal{D}(S) = \max_{\sigma, \sigma' \in S} d(\sigma, \sigma')$.

Lemma 9.1. $\frac{1}{2} \cdot \mathcal{D}(\mathcal{B}_t(\pi)) \leqslant k(t, \pi) \leqslant \mathcal{D}(\mathcal{B}_t(\pi)) \leqslant 2t$.

Proof. Let $\widehat{\sigma} = \text{MINIMAX}(\mathcal{B}_t(\pi))$. For rankings $\sigma, \sigma' \in \mathcal{B}_t(\pi)$, we have $d(\sigma, \widehat{\sigma}), d(\sigma', \widehat{\sigma}) \leqslant k(t, \pi)$ by definition of $\widehat{\sigma}$. By the triangle inequality, $d(\sigma, \sigma') \leqslant 2k(t, \pi)$ for all $\sigma, \sigma' \in \mathcal{B}_t(\pi)$. Thus, $\mathcal{D}(\mathcal{B}_t(\pi)) \leqslant 2k(t, \pi)$.

Next, the maximum distance of $\sigma \in \mathcal{B}_t(\pi)$ from all rankings in $\mathcal{B}_t(\pi)$ is at most $\mathcal{D}(\mathcal{B}_t(\pi))$. Hence, the minimax distance $k(t, \pi) = k(\mathcal{B}_t(\pi))$ cannot be greater than $\mathcal{D}(\mathcal{B}_t(\pi))$.

Finally, let $\pi = \{\sigma_1, \ldots, \sigma_n\}$. For rankings $\sigma, \sigma' \in \mathcal{B}_t(\pi)$, the triangle inequality implies $d(\sigma, \sigma') \leqslant d(\sigma, \sigma_i) + d(\sigma_i, \sigma')$ for every $i \in \{1, \ldots, n\}$. Averaging over these inequalities, we get $d(\sigma, \sigma') \leqslant t + t = 2t$, for all $\sigma, \sigma' \in \mathcal{B}_t(\pi)$. Thus, we have $\mathcal{D}(\mathcal{B}_t(\pi)) \leqslant 2t$, as required. ■

Lemma 9.1 implies that $k(t) = \max_{\pi \in \mathcal{L}(A)^n} k(t, \pi) \leqslant 2t$ for all distance metrics and $t > 0$. In words:

[3] Multiple rankings indicate a tie that can be broken arbitrarily.

Theorem 9.1. *Given n noisy rankings at an average distance of at most t from an unknown true ranking σ^* according to a distance metric d, it is always possible to find a ranking at distance at most $2t$ from σ^* according to d.*

Importantly, the bound of Theorem 9.1 is independent of the number of votes n. Most statistical models of social choice restrict profiles in two ways: i) the average error should be low because the probability of generating high-error votes is typically low, and ii) the errors should be distributed almost evenly (in different directions from the ground truth), which is why aggregating the votes works well. These assumptions are mainly helpful when n is large, that is, performance may be poor for small n (see, e.g., Chapter 7). In contrast, our model restricts profiles only by making the first assumption (explicitly), allowing voting rules to perform well as long as the votes are accurate on average, independently of the number of votes n.

We also remark that Theorem 9.1 admits a simple proof, but the bound is nontrivial: while the average error of the profile is at most t (hence, the profile contains a ranking with error at most t), it is generally impossible to pinpoint a single ranking within the profile that has error at most $2t$ with respect to the ground truth in the worst-case (i.e., with respect to every possible ground truth in $\mathcal{B}_t(\pi)$). That said, it can be shown that there exists a ranking in the profile that always has distance at most $3t$ from the ground truth. Further, one can pick such a ranking in polynomial time, which stands in sharp contrast to the usual hardness of finding the optimal ranking (see the discussion on the computational complexity of our approach in Section 9.7).

Theorem 9.2. *Given n noisy rankings at an average distance of at most t from an unknown true ranking σ^* according to a distance metric d, it is always possible to pick, in polynomial time, one of the n given rankings that has distance at most $3t$ from σ^* according to d.*

Proof. Consider a profile π consisting of n rankings such that $d(\sigma^*, \pi) \leqslant t$. Let $x = \min_{\sigma \in \mathcal{L}(A)} d(\sigma, \pi)$ be the minimum distance any ranking has from the profile. Then, $x \leqslant d(\sigma^*, \pi) \leqslant t$. Let $\widehat{\sigma} = \arg\min_{\sigma \in \pi} d(\sigma, \pi)$ be the ranking in π which minimizes the distance from π among all rankings in π. An easy-to-verify folklore theorem says that $d(\widehat{\sigma}, \pi) \leqslant 2x$. To see this, assume that ranking τ has the minimum distance from the profile (i.e., $d(\tau, \pi) = x$). Now, the average distance of all rankings in π from π is

$$\begin{aligned}\frac{1}{n}\sum_{\sigma\in\pi} d(\sigma,\pi) = \frac{1}{n^2}\sum_{\sigma\in\pi}\sum_{\sigma'\in\pi} d(\sigma,\sigma') &\leqslant \frac{1}{n^2}\sum_{\sigma\in\pi}\sum_{\sigma'\in\pi}\left(d(\tau,\sigma)+d(\tau,\sigma')\right)\\ &= \frac{2}{n}\sum_{\sigma\in\pi} d(\tau,\sigma) = 2x \leqslant 2t,\end{aligned}$$

where the second transition uses the triangle inequality. Now, $\widehat{\sigma}$ has the smallest distance from π among all rankings in π, which cannot be greater than the average distance $(1/n)\sum_{\sigma\in\pi} d(\sigma, \pi)$. Hence, $d(\widehat{\sigma}, \pi) \leqslant 2t$. Finally,

$$d(\widehat{\sigma},\sigma^*) \leqslant \frac{1}{n}\sum_{\sigma\in\pi}\left(d(\widehat{\sigma},\sigma)+d(\sigma,\sigma^*)\right) \leqslant 2t + t = 3t,$$

where the first transition uses the triangle inequality and the second transition uses the fact that $d(\widehat{\sigma}, \pi) \leqslant 2t$ and $d(\pi, \sigma^*) \leqslant t$. It is easy to see that $\widehat{\sigma}$ can be computed in $O(n^2)$ time. ■

9.3.2 Lower Bounds

The upper bound of $2t$ (Theorem 9.1) is intuitively loose — we cannot expect it to be tight for every distance metric. However, we can complement it with a lower bound of (roughly speaking) $t/2$ for all distance metrics. Formally, let $d^{\downarrow}(r)$ denote the greatest feasible distance under distance metric d that is less than or equal to r. Next, we prove a lower bound of $d^{\downarrow}(t)/2$.

Theorem 9.3. *For a distance metric d, $k(t) \geqslant d^{\downarrow}(t)/2$.*

Proof. If $d^{\downarrow}(t) = 0$, then the result trivially holds. Assume $d^{\downarrow}(t) > 0$. Let σ and σ' be two rankings at distance $d^{\downarrow}(t)$. Consider profile π consisting of only a single instance of ranking σ. Then, $\sigma' \in \mathcal{B}_t(\pi)$. Hence, $\mathcal{D}(\mathcal{B}_t(\pi)) \geqslant d^{\downarrow}(t)$. Now, it follows from Lemma 9.1 that $k(t) \geqslant \mathcal{D}(\mathcal{B}_t(\pi))/2 \geqslant d^{\downarrow}(t)/2$. ■

Recall that Theorem 9.1 shows that $k(t) \leqslant 2t$. However, $k(t)$ is the minimax distance under some profile, and hence must be a feasible distance under d. Thus, Theorem 9.1 actually implies a possibly better upper bound of $d^{\downarrow}(2t)$. Together with Theorem 9.3, this implies $d^{\downarrow}(t)/2 \leqslant k(t) \leqslant d^{\downarrow}(2t)$. Next, we show that imposing a mild assumption on the distance metric allows us to improve the lower bound by a factor of 2, thus reducing the gap between the lower and upper bounds.

Theorem 9.4. *For a neutral distance metric d, $k(t) \geqslant d^{\downarrow}(t)$.*

Proof. For a ranking $\sigma \in \mathcal{L}(A)$ and $r \geqslant 0$, let $\mathcal{B}_r(\sigma)$ denote the set of rankings at distance at most r from σ. Neutrality of the distance metric d implies $|\mathcal{B}_r(\sigma)| = |\mathcal{B}_r(\sigma')|$ for all $\sigma, \sigma' \in \mathcal{L}(A)$ and $r \geqslant 0$. In particular, $d^{\downarrow}(t)$ being a feasible distance under d implies that for every $\sigma \in \mathcal{L}(A)$, there exists some ranking at distance exactly $d^{\downarrow}(t)$ from σ.

Fix $\sigma \in \mathcal{L}(A)$. Consider the profile π consisting of n instances of σ. It holds that $\mathcal{B}_t(\pi) = \mathcal{B}_t(\sigma)$. We want to show that the minimax distance $k(\mathcal{B}_t(\sigma)) \geqslant d^{\downarrow}(t)$. Suppose for contradiction that there exists some $\sigma' \in \mathcal{L}(A)$ such that all rankings in $\mathcal{B}_t(\sigma)$ are at distance at most t' from σ', i.e., $\mathcal{B}_t(\sigma) \subseteq \mathcal{B}_{t'}(\sigma')$, with $t' < d^{\downarrow}(t)$. Since there exists some ranking at distance $d^{\downarrow}(t) > t'$ from σ', we have $\mathcal{B}_t(\sigma) \subseteq \mathcal{B}_{t'}(\sigma') \subsetneq \mathcal{B}_t(\sigma')$, which is a contradiction because $|\mathcal{B}_t(\sigma)| = |\mathcal{B}_t(\sigma')|$. Therefore, $k(t) \geqslant k(t, \pi) \geqslant d^{\downarrow}(t)$. ■

The bound of Theorem 9.4 holds for all $n, m > 0$ and all $t \in [0, D]$, where D is the maximum possible distance under d. It can be checked easily that the bound is tight given the neutrality assumption, which is an extremely mild — and in fact, a highly desirable — assumption for distance metrics over rankings.

Theorem 9.4 improves the bounds on $k(t)$ to $d^{\downarrow}(t) \leqslant k(t) \leqslant d^{\downarrow}(2t)$ for a variety of distance metrics d. However, for the four special distance metrics considered in this chapter, the next result, which is our main theoretical result, closes this gap by establishing a tight lower bound of $d^{\downarrow}(2t)$, for a wide range of values of n and t.

Theorem 9.5. *If $d \in \{d_{KT}, d_{FR}, d_{MD}, d_{CY}\}$, and the maximum distance allowed by the metric is $D \in \Theta(m^\alpha)$, then there exists $T \in \Theta(m^\alpha)$ such that:*

1. *For all $t \leqslant T$ and even n, we have $k(t) \geqslant d^{\downarrow}(2t)$.*
2. *For all $L \geqslant 2$, $t \leqslant T$ with $\{2t\} \in (1/L, 1-1/L)$, and odd $n \geqslant \Theta(L \cdot D)$, we have $k(t) \geqslant d^{\downarrow}(2t)$. Here, $\{x\} = x - \lfloor x \rfloor$ denotes the fractional part of $x \in \mathbb{R}$.*

The impossibility result of Theorem 9.5 is weaker for odd values of n (in particular, covering more values of t requires larger n), which is reminiscent of the fact that repetition (error-correcting) codes achieve greater efficiency with an odd number of repetitions; this is not merely a coincidence. Indeed, an extra repetition allows differentiating between tied possibilities for the ground truth; likewise, an extra vote in the profile prevents us from constructing a symmetric profile that admits a diverse set of possible ground truths.

Proof of Theorem 9.5. We denote $\{1, \ldots, r\}$ by $[r]$ in this proof. We use $\sigma(a)$ to denote the rank (position) of alternative a in ranking σ. First, we prove the case of even n for all four distance metrics. We later provide a generic argument to prove the case of large odd n. First, we need a simple observation.

Observation 9.1. *If $\binom{r}{2} \leqslant \lfloor 2t \rfloor$ and $t \geqslant 0.5$, then $r \leqslant 4\sqrt{t}$.*

Proof. Note that $(r-1)^2 \leqslant r \cdot (r-1) \leqslant 2 \cdot \lfloor 2t \rfloor \leqslant 4t$. Hence, $r \leqslant 2\sqrt{t} + 1$. We also have $t \geqslant 0.5$, i.e., $1 \leqslant 2t$. This implies $1 \leqslant \sqrt{2t}$. Thus, we have $r \leqslant 2\sqrt{t} + \sqrt{2t} = (2+\sqrt{2})\sqrt{t} \leqslant 4\sqrt{t}$. ■

The Kendall Tau Distance: Let d be the Kendall tau distance; thus, $D = \binom{m}{2}$ and $\alpha = 2$. Let n be even. For a ranking $\tau \in \mathcal{L}(A)$, let τ_{rev} be its reverse. Assume $t = (1/2) \cdot \binom{m}{2}$, and fix a ranking $\sigma \in \mathcal{L}(A)$. Every ranking must agree with exactly one of σ and σ_{rev} on a given pair of alternatives. Hence, every $\rho \in \mathcal{L}(A)$ satisfies $d(\rho, \sigma) + d(\rho, \sigma_{\text{rev}}) = \binom{m}{2}$. Consider the profile π consisting of $n/2$ instances of σ and $n/2$ instances of σ_{rev}. Then, the average distance of every ranking from rankings in π would be exactly t, i.e., $\mathcal{B}_t(\pi) = \mathcal{L}(A)$. It is easy to check that $k(\mathcal{L}(A)) = \binom{m}{2} = 2t = d^{\downarrow}(2t)$ because every ranking has its reverse ranking in $\mathcal{L}(A)$ at distance exactly $2t$.

Now, let us extend the proof to $t \leqslant (m/12)^2$. If $t < 0.5$, then $d^{\downarrow}_{KT}(2t) = 0$, which is a trivial lower bound. Hence, assume $t \geqslant 0.5$. Thus, $d^{\downarrow}(2t) = \lfloor 2t \rfloor$. We use Fermat's Polygonal Number Theorem (see, e.g., [110]). A special case of this remarkable theorem states that every natural number can be expressed as the sum of at most three "triangular" numbers, i.e., numbers of the form $\binom{k}{2}$. Let $\lfloor 2t \rfloor = \sum_{i=1}^{3} \binom{m_i}{2}$. From Observation 9.1, it follows that $0 \leqslant m_i \leqslant 4\sqrt{t}$ for all $i \in \{1,2,3\}$. Hence, $\sum_{i=1}^{3} m_i \leqslant 12\sqrt{t} \leqslant m$.

Partition the set of alternatives A into four disjoint groups A_1, A_2, A_3, and A_4 such that $|A_i| = m_i$ for $i \in \{1,2,3\}$, and $|A_4| = m - \sum_{i=1}^{3} m_i$. Let σ^{A_4} be an arbitrary ranking of the alternatives in A_4; consider the partial order $\mathcal{P}_A = A_1 \succ A_2 \succ A_3 \succ \sigma^{A_4}$ over alternatives in A. Note that a ranking ρ is an extension of $\mathcal{P}_A$ iff it ranks all alternatives in A_i before any alternative in A_{i+1} for $i \in \{1,2,3\}$, and ranks alternatives in A_4 according to σ^{A_4}. Choose arbitrary $\sigma^{A_i} \in \mathcal{L}(A_i)$ for $i \in \{1, 2, 3\}$ and define

$$\sigma = \sigma^{A_1} \succ \sigma^{A_2} \succ \sigma^{A_3} \succ \sigma^{A_4},$$
$$\sigma' = \sigma^{A_1}_{\text{rev}} \succ \sigma^{A_2}_{\text{rev}} \succ \sigma^{A_3}_{\text{rev}} \succ \sigma^{A_4}.$$

Note that both σ and σ' are extensions of $\mathcal{P}_A$. Once again, take the profile π consisting of $n/2$ instances of σ and $n/2$ instances of σ'. It is easy to check that a ranking disagrees with exactly one of σ and σ' on every pair of alternatives that belong to the same group in $\{A_1, A_2, A_3\}$. Hence, every ranking $\rho \in \mathcal{L}(A)$ satisfies

$$d(\rho, \sigma) + d(\rho, \sigma') \geqslant \sum_{i=1}^{3} \binom{m_i}{2} = \lfloor 2t \rfloor. \tag{9.1}$$

Clearly an equality is achieved in Equation (9.1) if and only if ρ is an extension of $\mathcal{P}_A$. Thus, every extension of $\mathcal{P}_A$ has an average distance of $\lfloor 2t \rfloor / 2 \leqslant t$ from π. Every ranking ρ that is not an extension of $\mathcal{P}_A$ achieves a strict inequality in Equation (9.1); thus, $d(\rho, \pi) \geqslant (\lfloor 2t \rfloor + 1)/2 > t$. Hence, $\mathcal{B}_t(\pi)$ is the set of extensions of $\mathcal{P}_A$.

Given a ranking $\rho \in \mathcal{L}(A)$, consider the ranking in $\mathcal{B}_t(\pi)$ that reverses the partial orders over A_1, A_2, and A_3 induced by ρ. The distance of this ranking from ρ would be at least $\sum_{i=1}^{3} \binom{m_i}{2} = \lfloor 2t \rfloor$, implying $k(\mathcal{B}_t(\pi)) \geqslant \lfloor 2t \rfloor$. (In fact, it can be checked that $k(\mathcal{B}_t(\pi)) = \mathcal{D}(\mathcal{B}_t(\pi)) = \lfloor 2t \rfloor$.)

We now proceed to prove the case of an even number of agents for the other three distance metrics. First, if M is the minimum distance between two distinct rankings under a distance metric d and $t < M/2$, then we have $d^{\downarrow}(2t) = 0$, which is a trivial lower bound. Hence, we assume $t \geqslant M/2$.

The Footrule Distance: Let d_{FR} denote the footrule distance; recall that given $\sigma, \sigma' \in \mathcal{L}(A)$, $d_{FR}(\sigma, \sigma') = \sum_{a \in A} |\sigma(a) - \sigma'(a)|$. The proof is along the same lines as the proof for the Kendall tau distance, but uses a few additional clever ideas. It is known that the maximum footrule distance between two rankings over m alternatives is $D = \lfloor m^2/2 \rfloor$, and is achieved by two rankings that are reverse of each other [75]. Hence, we have $\alpha = 2$; thus, we wish to find $T \in \Theta(m^2)$ for which the claim will hold. Formally writing the distance between a ranking and its reverse, we get

$$d_{FR}(\sigma, \sigma_{\text{rev}}) = \sum_{i=1}^{m} |m + 1 - 2i| = \left\lfloor \frac{m^2}{2} \right\rfloor. \tag{9.2}$$

Observation 9.2. *The footrule distance between two rankings is always an even integer.*

Proof. Take rankings $\sigma, \tau \in \mathcal{L}(A)$. Note that $d_{FR}(\sigma, \tau) = \sum_{a \in A} |\sigma(a) - \tau(a)|$. Now, $|\sigma(a) - \tau(a)|$ is odd if and only if the positions of a in σ and τ have different parity. Since the number of odd (as well as even) positions is identical in σ and τ, the number of alternatives that leave an even position in σ to go to an odd position in τ equals the number of alternatives that leave an odd position in σ to go to an even position in τ.

Thus, the number of alternatives for which the parity of the position changes is even. Equivalently, the number of odd terms in the sum defining the footrule distance is even. Hence, the footrule distance is an even integer. ■

Hence, Equation (9.2) implies that $d^{\downarrow}_{FR}(2t)$ equals $\lfloor 2t \rfloor$ if $\lfloor 2t \rfloor$ is even, and equals $\lfloor 2t \rfloor - 1$ otherwise. Let $r = d^{\downarrow}_{FR}(2t)$. Hence, r is an even integer. We prove the result for $t \leqslant (m/8)^2$. In this case, we invoke the 4-gonal special case of Fermat's Polygonal Number Theorem (instead of the 3-gonal case invoked in the proof for the Kendall tau distance): *Every positive integer can be written as the sum of at most four squares.* Let $r/2 = m_1^2 + m_2^2 + m_3^2 + m_4^2$. Hence,

$$r = \frac{(2m_1)^2}{2} + \frac{(2m_2)^2}{2} + \frac{(2m_3)^2}{2} + \frac{(2m_4)^2}{2}. \tag{9.3}$$

It is easy to check that $m_i \leqslant \sqrt{r/2}$ for $i \in [4]$. Thus, $\sum_{i=1}^{4} 2m_i \leqslant 8\sqrt{r/2} \leqslant 8\sqrt{t} \leqslant m$. Let us partition the set of alternatives A into $\{A_i\}_{i \in [5]}$ such that $|A_i| = 2m_i$ for $i \in [4]$ and $|A_5| = m_5 = m - \sum_{i=1}^{4} 2m_i$.

Fix $\sigma^{A_5} \in \mathcal{L}(A_5)$ and consider the partial order $\mathcal{P}_A = A_1 \succ A_2 \succ A_3 \succ A_4 \succ \sigma^{A_5}$. Choose arbitrary $\sigma^{A_i} \in \mathcal{L}(A_i)$ for $i \in [4]$, and let

$$\begin{aligned} \sigma &= (\sigma^{A_1} \succ \sigma^{A_2} \succ \sigma^{A_3} \succ \sigma^{A_4} \succ \sigma^{A_5}), \\ \sigma' &= (\sigma^{A_1}_{\text{rev}} \succ \sigma^{A_2}_{\text{rev}} \succ \sigma^{A_3}_{\text{rev}} \succ \sigma^{A_4}_{\text{rev}} \succ \sigma^{A_5}). \end{aligned}$$

Note that both σ and σ' are extensions of $\mathcal{P}_A$. Consider the profile π consisting of $n/2$ instances of σ and σ' each. Unlike the Kendall tau distance, $\mathcal{B}_t(\pi)$ is not the set of extensions of $\mathcal{P}_A$. Still, we show that it satisfies $k(\mathcal{B}_t(\pi)) = \mathcal{D}(\mathcal{B}_t(\pi)) = d^{\downarrow}_{FR}(2t) = r$.

Denote by a_i^j the alternative ranked j in σ^{A_i}. Take a ranking $\rho \in \mathcal{L}(A)$. Consider $d_{FR}(\rho, \sigma) + d_{FR}(\rho, \sigma')$. We have the following inequalities regarding the sum of displacement of different alternatives between ρ and σ, and between ρ and σ'. For $i \in [4]$ and $j \in [2m_i]$,

$$\left|\rho(a_i^j) - \sigma(a_i^j)\right| + \left|\rho(a_i^j) - \sigma'(a_i^j)\right| \geqslant \left|\sigma(a_i^j) - \sigma'(a_i^j)\right| = |j - (2m_i + 1 - j)|. \tag{9.4}$$

Summing all the inequalities, we get

$$d_{FR}(\rho, \sigma) + d_{FR}(\rho, \sigma') \geqslant \sum_{i=1}^{4} \sum_{j=1}^{2m_i} |2j - 2m_i - 1| = \sum_{i=1}^{4} \frac{(2m_i)^2}{2} = r, \tag{9.5}$$

where the second transition follows from Equation (9.2), and the third transition follows from Equation (9.3).

First, we show that $\rho \in \mathcal{B}_t(\pi)$ only if equality in Equation (9.5) holds. To see why, note that the footrule distance is always even and $r = d^{\downarrow}_{FR}(2t) \geqslant \lfloor 2t \rfloor - 1$. Hence, if equality is not achieved, then $d_{FR}(\rho, \sigma) + d_{FR}(\rho, \sigma') \geqslant r + 2 \geqslant \lfloor 2t \rfloor - 1 + 1 > 2t$. Hence, the average distance of ρ from votes in π would be greater than t.

On the contrary, if equality is indeed achieved in Equation (9.5), then the average distance of ρ from votes in π is $r/2 \leqslant t$. Hence, we have established that $\mathcal{B}_t(\pi)$ is the set of rankings ρ for which equality is achieved in Equation (9.5).

For ρ to achieve equality in Equation (9.5), it must achieve equality in Equation (9.4) for every $i \in [4]$ and $j \in [2m_i]$, and it must agree with both σ and σ' on the positions of alternatives in A_5 (i.e., σ^{A_5} must be a suffix of ρ). For the former to hold, the position of a_i^j in ρ must be between $\sigma(a_i^j)$ and $\sigma'(a_i^j) = \sigma(a_i^{2m_i+1-j})$ (both inclusive), for every $i \in [4]$ and $j \in [2m_i]$.

We claim that the set of rankings satisfying these conditions are characterized as follows.

$$\begin{aligned}\mathcal{B}_t(\pi) = \Big\{\rho \in \mathcal{L}(A) \;\Big|\; & \{\rho(a_i^j), \rho(a_i^{2m_i+1-j})\} = \{\sigma(a_i^j), \sigma(a_i^{2m_i+1-j})\} \\ & \text{for } i \in [4], j \in [2m_i], \text{ and} \\ & \rho(a_5^j) = \sigma(a_5^j) = \sigma'(a_5^j) \text{ for } j \in [m_5]\Big\}. \end{aligned} \tag{9.6}$$

Note that instead of $\rho(a_i^j)$ and $\rho(a_i^{2m_i+1-j})$ both being in the interval $[\sigma(a_i^j), \sigma(a_i^{2m_i+1-j})]$, we are claiming that they must be the two endpoints. First, consider the middle alternatives in each A_i ($i \in [4]$), namely $a_i^{m_i}$ and $a_i^{m_i+1}$. Both must be placed between $\sigma(a_i^{m_i}) = \sigma'(a_i^{m_i+1})$ and $\sigma(a_i^{m_i+1}) = \sigma'(a_i^{m_i})$; but these two numbers differ by exactly 1. Hence,

$$\left\{\rho(a_i^{m_i}), \rho(a_i^{m_i+1})\right\} = \left\{\sigma(a_i^{m_i}), \sigma(a_i^{m_i+1})\right\}.$$

Consider the two adjacent alternatives, namely $a_i^{m_i-1}$ and $a_i^{m_i+2}$. Given that the middle alternatives $a_i^{m_i}$ and $a_i^{m_i+1}$ occupy their respective positions in σ or σ', the only positions available to ρ for placing the two adjacent alternatives are the endpoints of their common feasible interval $[\sigma(a_i^{m_i-1}), \sigma(a_i^{m_i+2})]$. Continuing this argument, each pair of alternatives $(a_i^j, a_i^{2m_i+1-j})$ must occupy the two positions $\{\sigma(a_i^j), \sigma(a_i^{2m_i+1-j})\}$ for every $i \in [4]$ and $j \in [m_i]$.

That is, ρ can either keep the alternatives a_i^j and $a_i^{2m_i+1-j}$ as they are in σ, or place them according to σ' (equivalently, swapping them in σ) for every $i \in [4]$ and $j \in [2m_i]$. Note that these choices are independent of each other. We established that a ranking ρ is in $\mathcal{B}_t(\pi)$ only if it is obtained in this manner and has σ^{A_5} as its suffix.

Further, it can be seen that each of these choices (keeping or swapping the pair in σ) maintain $d_{FR}(\rho, \sigma) + d_{FR}(\rho, \sigma')$ invariant. Hence, all such rankings ρ satisfy $d_{FR}(\rho, \sigma) + d_{FR}(\rho, \sigma') = r$, and thus belong to $\mathcal{B}_t(\pi)$. This reaffirms our original claim that $\mathcal{B}_t(\pi)$ is given by Equation (9.6).

In summary, all rankings in $\mathcal{B}_t(\pi)$ can be obtained by taking σ, and arbitrarily choosing whether to swap the pair of alternatives a_i^j and $a_i^{2m_i+1-j}$ for each $i \in [4]$ and $j \in [2m_i]$.

Note that $\sigma, \sigma' \in \mathcal{B}_t(\pi)$ and $d_{FR}(\sigma, \sigma') = r$ (this distance is given by the summation in Equation (9.5)). Hence, $\mathcal{D}(\mathcal{B}_t(\pi)) \geqslant r$. Now, we prove that its minimax distance is at least r as well. Take a ranking $\rho \in \mathcal{L}(A)$. We need to show that there exists some $\tau \in \mathcal{B}_t(\pi)$ such that $d_{FR}(\rho, \tau) \geqslant r$.

Consider alternatives a_i^j and $a_i^{2m_i+1-j}$ for $i \in [4]$ and $j \in [2m_i]$. We know that τ must satisfy $\{\tau(a_i^j), \tau(a_i^{2m_i+1-j})\} = \{\sigma(a_i^j), \sigma(a_i^{2m_i+1-j})\}$ in order to belong to $\mathcal{B}_t(\pi)$. This allows two possible ways for placing the pair of alternatives. Let τ pick the optimal positions that maximize

$$\tau_{i,j}(\rho) = |\tau(a_i^j) - \rho(a_i^j)| + |\tau(a_i^{2m_i+1-j}) - \rho(a_i^{2m_i+1-j})|.$$

That is, $\tau_{i,j}(\rho)$ should equal $M_{i,j}(\rho)$, which we define as

$$\max \Big\{ |\sigma(a_i^j) - \rho(a_i^j)| + |\sigma(a_i^{2m_i+1-j}) - \rho(a_i^{2m_i+1-j})|, \\ |\sigma(a_i^{2m_i+1-j}) - \rho(a_i^j)| + |\sigma(a_i^j) - \rho(a_i^{2m_i+1-j})| \Big\}.$$

Note that the choice for each pair of alternatives $(a_i^j, a_i^{2m_i+1-j})$ can be made independently of every other pair. Further, making the optimal choice for each pair guarantees that $d_{FR}(\rho, \tau)$ is at least

$$\sum_{i=1}^{4} \sum_{j=1}^{2m_i} \tau_{i,j}(\rho) = \sum_{i=1}^{4} \sum_{j=1}^{2m_i} M_{i,j}(\rho),$$

which we will now show to be at least r.

Algorithm 8 describes how to find the optimal ranking $\tau \in \mathcal{B}_t(\pi)$ mentioned above, which satisfies $\tau_{i,j}(\rho) = M_{i,j}(\rho)$ for every $i \in [4]$ and $j \in [2m_i]$. It starts with an arbitrary $\tau \in \mathcal{B}_t(\pi)$, and swaps every sub-optimally placed pair $(a_i^j, a_i^{2m_i+1-j})$ for $i \in [4]$ and $j \in [2m_i]$. In the algorithm, $\tau_{a \leftrightarrow b}$ denotes the ranking obtained by swapping alternatives a and b in τ.

Finally, we show that $d_{FR}(\rho, \tau) \geqslant r$. First, we establish the following lower bound on $M_{i,j}(\rho)$.

$$\begin{aligned} M_{i,j}(\rho) \geqslant & \frac{1}{2} \Big(|\sigma(a_i^j) - \rho(a_i^j)| + |\sigma(a_i^{2m_i+1-j}) - \rho(a_i^{2m_i+1-j})| \\ & + |\sigma(a_i^{2m_i+1-j}) - \rho(a_i^j)| + |\sigma(a_i^j) - \rho(a_i^{2m_i+1-j})| \Big) \\ \geqslant & |\sigma(a_i^{2m_i+1-j}) - \sigma(a_i^j)| \\ = & |2m_i + 1 - 2j|, \end{aligned}$$

where the first transition holds because the maximum of two terms is at least as much as their average, and the second transition uses the triangle inequality on appropriately paired terms. Now, we have

$$d_{FR}(\tau, \rho) \geqslant \sum_{i=1}^{4} \sum_{j=1}^{2m_i} M_{i,j}(\rho) \geqslant \sum_{i=1}^{4} \sum_{j=1}^{2m_i} |2m_i + 1 - 2j| = \sum_{i=1}^{4} \frac{(2m_i)^2}{2} = r,$$

ALGORITHM 8: Finds a ranking in $\mathcal{B}_t(\pi)$ at footrule distance at least $\lfloor 2t \rfloor$ from a given ranking.

Data: Ranking $\rho \in \mathcal{L}(A)$
Result: Ranking $\tau \in \mathcal{B}_t(\pi)$ such that $d_{FR}(\tau, \rho) \geqslant \lfloor 2t \rfloor$

```
τ ← an arbitrary ranking from B_t(π);
for i ∈ [4] do
    for j ∈ [2m_i] do
        d_i^j ← |ρ(a_i^j) − τ(a_i^j)|;
        d_i^{2m_i+1−j} ← |ρ(a_i^{2m_i+1−j}) − τ(a_i^{2m_i+1−j})|;
        if d_i^j + d_i^{2m_i+1−j} < M_{i,j}(ρ) then
            τ ← τ_{a_i^j ↔ a_i^{2m_i+1−j}};
        end
    end
end
return τ;
```

where the third transition holds due to Equation (9.2), and the fourth transition holds due to Equation (9.3). Hence, the minimax distance of $\mathcal{B}_t(\pi)$ is at least $r = d^{\downarrow}_{FR}(2t)$, as required.

The Cayley Distance: Next, let d_{CY} denote the Cayley distance. Recall that $d_{CY}(\sigma, \tau)$ equals the minimal number of swaps (of possibly non-adjacent alternatives) required in order to transform σ to τ. It is easy to check that the maximum Cayley distance is $D = m - 1$; hence, it has $\alpha = 1$. We prove the result for $t \leqslant m/4$. Note that $d^{\downarrow}_{CY}(2t) = \lfloor 2t \rfloor$. Define rankings $\sigma, \sigma' \in \mathcal{L}(A)$ as follows.

$$\sigma = (\underbrace{a_1 \succ \ldots \succ a_{2\lfloor 2t \rfloor}} \succ a_{2\lfloor 2t \rfloor + 1} \succ \ldots \succ a_m),$$
$$\sigma' = (\underbrace{a_{2\lfloor 2t \rfloor} \succ \ldots \succ a_1} \succ a_{2\lfloor 2t \rfloor + 1} \succ \ldots \succ a_m).$$

Let profile π consist of $n/2$ instances of σ and σ' each. We claim that $\mathcal{B}_t(\pi)$ has the following structure, which is very similar to the ball for the footrule distance.

$$\mathcal{B}_t(\pi) = \Big\{ \rho \in \mathcal{L}(A) \;\Big|\; \{\rho(a_i), \rho(a_{2\lfloor 2t \rfloor + 1 - i})\} = \{i, 2\lfloor 2t \rfloor + 1 - i\} \text{ for } i \in [\lfloor 2t \rfloor], \text{ and } \rho(a_i) = i \text{ for } i > 2\lfloor 2t \rfloor \Big\}. \tag{9.7}$$

First, we observe the following simple fact: If rankings τ and ρ mismatch (i.e., place different alternatives) in r different positions, then $d_{CY}(\tau, \rho) \geqslant r/2$. Indeed, consider the number of swaps required to convert τ into ρ. Since each swap can make τ and ρ consistent in at most two more positions, it would take at least $r/2$ swaps to convert τ into ρ, i.e., $d_{CY}(\tau, \rho) \geqslant r/2$.

Now, note that σ and σ' mismatch in each of first $2\lfloor 2t \rfloor$ positions. Hence, every ranking $\rho \in \mathcal{L}(A)$ must mismatch with at least one of σ and σ' in each of first $2\lfloor 2t \rfloor$

positions. Together with the previous observation, this implies

$$d_{CY}(\rho,\sigma)+d_{CY}(\rho,\sigma')\geqslant\lfloor 2t\rfloor\,. \tag{9.8}$$

Every ranking ρ that achieves equality in Equation (9.8) is clearly in $\mathcal{B}_t(\pi)$ because its average distance from the votes in π is $\lfloor 2t\rfloor/2\leqslant t$. Further, every ranking ρ that achieves a strict inequality in Equation (9.8) is outside $\mathcal{B}_t(\pi)$ because its average distance from the votes in π is at least $(\lfloor 2t\rfloor+1)/2>t$. Hence, $\mathcal{B}_t(\pi)$ consists of rankings that satisfy $d_{CY}(\rho,\sigma)+d_{CY}(\rho,\sigma')=\lfloor 2t\rfloor$.

Now, any ranking ρ satisfying equality in Equation (9.8) must be consistent with exactly one of σ and σ' in each of first $2\lfloor 2t\rfloor$ positions, and with both σ and σ' in the later positions. The former condition implies that for every $i\in\lfloor 2t\rfloor$, ρ must place the pair of alternatives $(a_i,a_{2\lfloor 2t\rfloor+1-i})$ in positions i and $2\lfloor 2t\rfloor+1-i$, either according to σ or according to σ'. This confirms our claim that $\mathcal{B}_t(\pi)$ is given by Equation (9.7).

We now show that $k(\mathcal{B}_t(\pi))\geqslant\lfloor 2t\rfloor$. Take a ranking $\rho\in\mathcal{L}(A)$. We construct a ranking $\tau\in\mathcal{B}_t(\pi)$ such that τ mismatches with ρ in each of first $2\lfloor 2t\rfloor$ positions. Together with our observation that the Cayley distance is at least half of the number of positional mismatches, this would imply that the minimax distance of $\mathcal{B}_t(\pi)$ is at least $\lfloor 2t\rfloor$, as required.

We construct τ by choosing the placement of the pair of alternatives $(a_i,a_{2\lfloor 2t\rfloor+1-i})$, independently for each $i\in\lfloor 2t\rfloor$, in a way that τ mismatches with ρ in positions i and $2\lfloor 2t\rfloor+1-i$ both. Let $\mathbb{I}(X)$ denote the indicator variable that is 1 if statement X holds, and 0 otherwise. Let $r=\mathbb{I}\left(\rho(a_i)=i\right)+\mathbb{I}\left(\rho(a_{2\lfloor 2t\rfloor+1-i})=2\lfloor 2t\rfloor+1-i\right)$. Consider the following three cases.

$r=0$: Set $\tau(a_i)=i$ and $\tau(a_{2\lfloor 2t\rfloor+1-i})=2\lfloor 2t\rfloor+1-i$.

$r=1$: Without loss of generality, assume $\rho(a_i)=i$. Set $\tau(a_i)=2\lfloor 2t\rfloor+1-i$ and $\tau(a_{2\lfloor 2t\rfloor+1-i})=i$.

$r=2$: Set $\tau(a_i)=2\lfloor 2t\rfloor+1-i$ and $\tau(a_{2\lfloor 2t\rfloor+1-i})=i$.

Finally, set $\tau(a_i)=i$ for all $i>2\lfloor 2t\rfloor$. This yields a ranking τ that is in $\mathcal{B}_t(\pi)$, and mismatches ρ in each of first $2\lfloor 2t\rfloor$ positions; hence, $d_{CY}(\rho,\tau)\geqslant\lfloor 2t\rfloor$, as required.

The Maximum Displacement Distance: Finally, let d_{MD} denote the maximum displacement distance. Note that it can be at most $D=m-1$; hence, it also has $\alpha=1$. However, this distance metric requires an entirely different technique than the ones used for previous distances. For example, taking any two rankings at maximum distance from each other does not work. We prove this result for $t\leqslant m/4$. Once again, note that $d^{\downarrow}_{MD}(2t)=\lfloor 2t\rfloor$.

Consider rankings σ and σ' defined as follows.

$$\sigma=(\underbrace{a_1\succ\ldots\succ a_{\lfloor 2t\rfloor}}\succ\underbrace{a_{\lfloor 2t\rfloor+1}\succ\ldots\succ a_{2\lfloor 2t\rfloor}}\succ a_{\text{rest}}),$$
$$\sigma'=(\underbrace{a_{\lfloor 2t\rfloor+1}\succ\ldots\succ a_{2\lfloor 2t\rfloor}}\succ\underbrace{a_1\succ\ldots\succ a_{\lfloor 2t\rfloor}}\succ a_{\text{rest}}),$$

where a_{rest} is shorthand for $a_{2\lfloor 2t \rfloor+1} \succ \ldots \succ a_m$. Note that the blocks of alternatives a_1 through $a_{\lfloor 2t \rfloor}$ and $a_{\lfloor 2t \rfloor+1}$ through $a_{2\lfloor 2t \rfloor}$ are shifted to each other's positions in the two rankings. Thus, each of a_1 through $a_{2\lfloor 2t \rfloor}$ have a displacement of exactly $\lfloor 2t \rfloor$ between the two rankings. Thus, $d_{MD}(\sigma, \sigma') = \lfloor 2t \rfloor$.

Consider the profile π consisting of $n/2$ instances of σ and σ' each. Clearly, σ and σ' have an average distance of $\lfloor 2t \rfloor / 2 \leqslant t$ from rankings in π. Hence, $\{\sigma, \sigma'\} \in \mathcal{B}_t(\pi)$. Surprisingly, in this case we can show that the minimax distance of $\mathcal{B}_t(\pi)$ without any additional information regarding the structure of $\mathcal{B}_t(\pi)$.

Take a ranking $\rho \in \mathcal{L}(A)$. The alternative placed first in ρ must be ranked at a position $\lfloor 2t \rfloor$ or below in at least one of σ and σ'. Hence, $\max(d_{MD}(\rho, \sigma), d_{MD}(\rho, \sigma')) \geqslant \lfloor 2t \rfloor$. Thus, there exists a ranking in $\mathcal{B}_t(\pi)$ at distance at least $\lfloor 2t \rfloor$ from ρ, i.e., the minimax distance of $\mathcal{B}_t(\pi)$ is at least $\lfloor 2t \rfloor$, as desired.

This completes the proof of the special case of even n for all four distance metrics. Now, consider the case of odd n.

Odd *n*: To extend the proof to odd values of n, we simply add one more instance of σ than σ'. The key insight is that with large n, the distance from the additional vote would have little effect on the average distance of a ranking from the profile. Thus, $\mathcal{B}_t(\pi)$ would be preserved, and the proof would follow.

Formally, let $L \geqslant 2$ and $t \in (1/L, 1 - 1/L)$. For the case of even n, the proofs for all four distance metrics proceeded as follows: Given the feasible distance $r = d^{\downarrow}(2t)$, we constructed two rankings σ and σ' at distance r from each other such that $\mathcal{B}_t(\pi)$ is the set of rankings at minimal total distance from the two rankings, i.e.,

$$\mathcal{B}_t(\pi) = \left\{ \rho \in \mathcal{L}(A) \mid d(\rho, \sigma) + d(\rho, \sigma') = r \right\}.$$

Let $n \geqslant 3$ be odd. Consider the profile π that has $(n-1)/2$ instances of σ and σ' each, and an additional instance of an arbitrary ranking. In our generic proof for all four distance metrics, we obtain conditions under which $\mathcal{B}_t(\pi) = \mathcal{B}_t(\pi')$ where π' is obtained by removing the arbitrary ranking from π (and hence has an even number of votes). We already proved that $k(\mathcal{B}_t(\pi')) \geqslant d^{\downarrow}(2t)$. Hence, obtaining $\mathcal{B}_t(\pi) = \mathcal{B}_t(\pi')$ would also show the lower bound $d^{\downarrow}(2t)$ for odd n.

In more detail, our objective is that every ranking ρ with $d(\rho, \sigma) + d(\rho, \sigma') = r$ (which may have a worst-case distance of D from the additional arbitrary ranking) should be in $\mathcal{B}_t(\pi)$, and every ranking ρ with $d(\rho, \sigma) + d(\rho, \sigma') > r$ should be outside $\mathcal{B}_t(\pi)$.

First, let $d \in \{d_{KT}, d_{CY}, d_{MD}\}$. If $d(\rho, \sigma) + d(\rho, \sigma') > r$, then $d(\rho, \sigma) + d(\rho, \sigma') \geqslant r + 1$. The total error incurred by rankings of distance r from π is $\frac{n-1}{2} \cdot r$, and a distance of D from the additional ranking. This means that

$$t \geqslant \frac{\frac{n-1}{2} \cdot r + D}{n}.$$

For rankings with an error greater than r to be outside $\mathcal{B}_t(\pi)$, we must have

$$t < \frac{\frac{n-1}{2} \cdot (r+1)}{n}.$$

Combining the inequalities, we obtain that

$$\begin{aligned}
&\frac{\frac{n-1}{2}\cdot r + D}{n} \leqslant t < \frac{\frac{n-1}{2}\cdot (r+1)}{n} \\
\Leftrightarrow &\frac{n-1}{2}\cdot r + D \leqslant n \cdot t < \frac{n-1}{2}\cdot (r+1) \\
\Leftrightarrow &r + \frac{2D}{n-1} \leqslant \frac{2n}{n-1}\cdot t < r+1 \\
\Leftrightarrow &r \leqslant 2t - \frac{2D-2t}{n-1} < r + \left(1 - \frac{2D}{n-1}\right). \qquad (9.9)
\end{aligned}$$

Choose $n \geqslant 2LD+1$. Then, $2D/(n-1) \leqslant 1/L < \{2t\}$. Note that

$$\left\lfloor 2t - \frac{2D-2t}{n-1} \right\rfloor = \left\lfloor \lfloor 2t \rfloor + \{2t\} - \frac{2D-2t}{n-1} \right\rfloor = \lfloor 2t \rfloor,$$

where the last equality holds because we showed $(2D-2t)/(n-1) < \{2t\}$.

In all three distance metrics considered thus far, we had $\lfloor 2t \rfloor = d^{\downarrow}(2t)$. Let $r = \lfloor 2t \rfloor$. We show that $r = \lfloor 2t \rfloor$ satisfies Equation (9.9), thus yielding $\mathcal{B}_t(\pi)$ with minimax distance at least $r = d^{\downarrow}(2t)$, as required. Note that

$$r \leqslant 2t - \frac{2D-2t}{n-1}$$

is satisfied by definition from Equation (9.9). We also have

$$\begin{aligned}
\left(2t - \frac{2D-2t}{n-1}\right) - \left(r + 1 - \frac{2D}{n-1}\right) &= 2t + \frac{2t}{n-1} - \lfloor 2t \rfloor - 1 \\
&= \{2t\} + \frac{2t}{n-1} - 1 \\
&< 1 - \frac{1}{L} + \frac{1}{L} - 1 = 0.
\end{aligned}$$

Hence, we have $k(t) \geqslant d^{\downarrow}(2t)$ for $n \geqslant 2LD+1$.

Next, consider the footrule distance. If $\lfloor 2t \rfloor$ is even (i.e., if $\lfloor 2t \rfloor = d^{\downarrow}(2t)$), then the above proof works because $r = \lfloor 2t \rfloor$ is a feasible distance. If $\lfloor 2t \rfloor$ is odd, then we must choose $r = \lfloor 2t \rfloor - 1$. However, we have an advantage: since the footrule distance is always even, every ranking ρ with $d(\rho,\sigma) + d(\rho,\sigma') > r$ must have $d(\rho,\sigma) + d(\rho,\sigma') \geqslant r+2$. Hence, we only need

$$\begin{aligned}
&\frac{\frac{n-1}{2}\cdot r + D}{n} \leqslant t < \frac{\frac{n-1}{2}\cdot (r+2)}{n} \\
\Leftrightarrow\ & r \leqslant 2t - \frac{2D-2t}{n-1} < r + \left(2 - \frac{2D}{n-1}\right). \qquad (9.10)
\end{aligned}$$

Note that $r = \lfloor 2t \rfloor - 1$ clearly satisfies the first inequality in Equation (9.10). For the second inequality, note that r decreased by 1 compared to earlier but $1 - 2D/(n-1)$ increased to $2 - 2D/(n-1)$ instead. Hence, the second inequality is still satisfied, and we get $\mathcal{B}_t(\pi)$ with minimax distance at least $r = \lfloor 2t \rfloor - 1 = d^{\downarrow}(2t)$, as required. ■

9.4 Approximations for Unknown Average Error

In the previous sections we derived the optimal rules when the upper bound t on the average error is given to us. In practice, the given bound may be inaccurate. We know that using an estimate $\widehat{t}$ that is still an upper bound ($\widehat{t} \geqslant t$) yields a ranking at distance at most $2\widehat{t}$ from the ground truth in the worst case. What happens if it turns out that $\widehat{t} < t$? We show that the output ranking is still at distance at most $4t$ from the ground truth in the worst case.

Theorem 9.6. *For a distance metric d, a profile π consisting of n noisy rankings at an average distance of at most t from the true ranking σ^*, and $\widehat{t} < t$, $d(\mathrm{OPT}^d(\widehat{t}, \pi), \sigma^*) \leqslant 4t$.*

To prove the theorem, we make a detour through minisum rules. For a distance metric d, let $\textsc{MiniSum}^d$, be the voting rule that always returns the ranking minimizing the sum of distances (equivalently, average distance) from the rankings in the given profile according to d. Two popular minisum rules are the Kemeny rule for the Kendall tau distance ($\textsc{MiniSum}^{d_{KT}}$) and the minisum rule for the footrule distance ($\textsc{MiniSum}^{d_{FR}}$), which approximates the Kemeny rule [79].[4] For a distance metric d (dropped from the superscripts), let $d(\pi, \sigma^*) \leqslant t$. We claim that the minisum ranking $\textsc{MiniSum}(\pi)$ is at distance at most $\min(2t, 2k(t, \pi))$ from σ^*. This is true because the minisum ranking and the true ranking are both in $\mathcal{B}_t(\pi)$, and Lemma 9.1 shows that its diameter is at most $\min(2t, 2k(t, \pi))$.

Returning to the theorem, if we provide an underestimate $\widehat{t}$ of the true worst-case average error t, then using Lemma 9.1,

$$
\begin{aligned}
&d\left(\textsc{MiniMax}(\mathcal{B}_{\widehat{t}}(\pi)), \textsc{MiniSum}(\pi)\right) \leqslant 2\widehat{t} \leqslant 2t,\\
&d\left(\textsc{MiniSum}(\pi), \sigma^*\right) \leqslant \mathcal{D}(\mathcal{B}_t(\pi)) \leqslant 2t.
\end{aligned}
$$

By the triangle inequality, $d\left(\textsc{MiniMax}(\mathcal{B}_{\widehat{t}}(\pi)), \sigma^*\right) \leqslant 4t$.

9.5 Experimental Results

We compare our worst-case optimal voting rules OPT^d against a plethora of voting rules used in the literature: plurality, Borda count, veto, the Kemeny rule, single transferable vote (STV), Copeland's rule, Bucklin's rule, the maximin rule, Slater's rule, Tideman's rule, and the modal ranking rule (for definitions see Chapter 7).

Our performance measure is the distance of the output ranking from the *actual* ground truth. In contrast, for a given d, OPT^d is designed to optimize the worst-case distance to *any possible* ground truth. Hence, crucially, OPT^d *is not guaranteed to outperform other rules in our experiments*.

[4]Minisum rules such as the Kemeny rule are also compelling because they often satisfy attractive social choice axioms. However, it is unclear whether such axioms contribute to the overall goal of effectively recovering the ground truth.

We use two real-world datasets containing ranked preferences in domains where ground truth rankings exist. Mao et al. [140] collected these datasets — *dots* and *puzzle* — via Amazon Mechanical Turk. For dataset *dots* (resp., *puzzle*), human workers were asked to rank four images that contain a different number of dots (resp., different states of an 8-Puzzle) according to the number of dots (resp., the distances of the states from the goal state). Each dataset has four different noise levels (i.e., levels of task difficulty), represented using a single noise parameter: for dots (resp., puzzle), higher noise corresponds to ranking images with a smaller difference between their number of dots (resp., ranking states that are all farther away from the goal state). Each dataset has 40 profiles with approximately 20 votes each, for each of the 4 noise levels. Points in our graphs are averaged over the 40 profiles in a single noise level of a dataset.

First, as a sanity check, we verified (Figure 9.1) that the noise parameter in the datasets positively correlates with our notion of noise — the average error in the profile, denoted t^* (averaged over all profiles in a noise level). Notably, the results from the two datasets are almost identical!

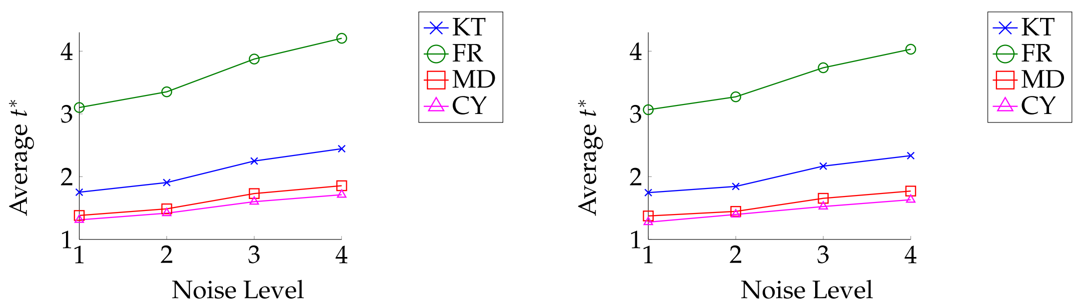

(a) Dots (b) Puzzle

Figure 9.1: Positive correlation of t^* with the noise parameter

Next, we compare OPT^d and MINISUM^d against the voting rules listed above, with distance d as the measure of error. We use the average error in a profile as the bound t given to OPT^d, i.e., we compute OPT$^d(t^*, \pi)$ on profile π where $t^* = d(\pi, \sigma^*)$. While this is somewhat optimistic, note that t^* may not be the (optimal) value of t that achieves the lowest error. Also, the experiments below show that a reasonable estimate of t^* also suffices.

Figures 9.2(a) and 9.2(b) show the results for the dots and puzzle datasets, respectively, under the Kendall tau distance. It can be seen that OPT$^{d_{KT}}$ (solid red line) significantly outperforms all other voting rules. The three other distance metrics considered in this chapter generate similar results; the corresponding graphs are presented in the appendix.

Finally, we test OPT^d in the more demanding setting where only an estimate $\widehat{t}$ of t^* is provided. To synchronize the results across different profiles, we use $r = (\widehat{t} - \text{MAD})/(t^* - \text{MAD})$, where MAD is the *minimum average distance* of any ranking from the votes in a profile, that is, the average distance of the ranking returned by MINISUM^d from the input votes. For all profiles, $r = 0$ implies $\widehat{t} = \text{MAD}$ (the smallest value that

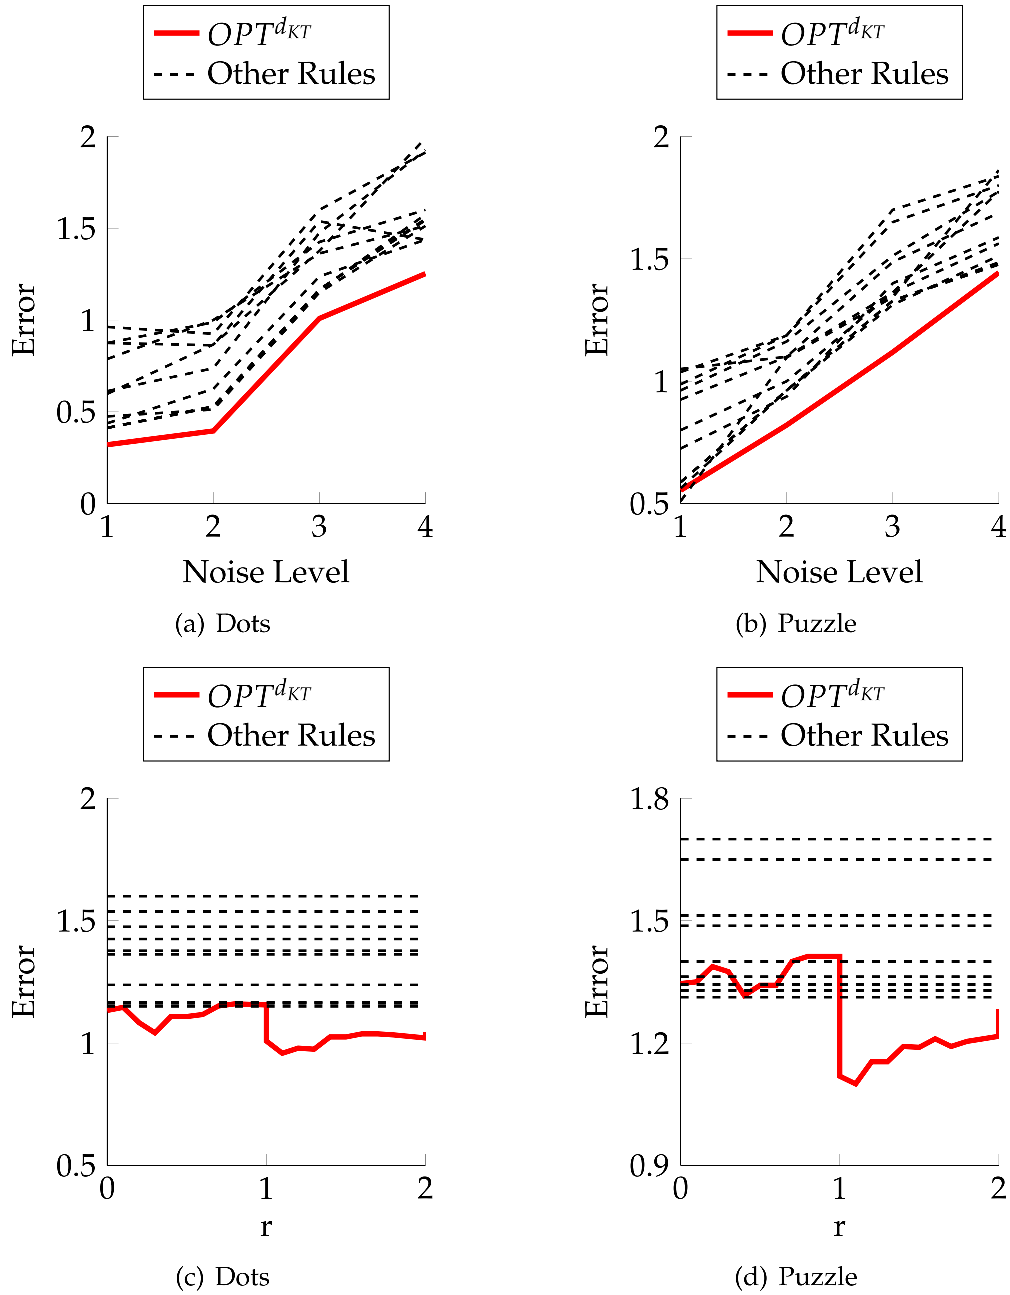

(a) Dots (b) Puzzle (c) Dots (d) Puzzle

Figure 9.2: Performance of different voting rules (Figures 9.2(a) and 9.2(b)), and of OPT with varying $\widehat{t}$ (Figures 9.2(c) and 9.2(d)).

admits a possible ground truth) and $r = 1$ implies $\widehat{t} = t^*$ (the true average error). In our experiments we use $r \in [0, 2]$; here, $\widehat{t}$ is an overestimate of t^* for $r \in (1, 2]$ (a valid upper bound on t^*), but an underestimate of t^* for $r \in [0, 1)$ (an invalid upper bound on t^*).

Figures 9.2(c) and 9.2(d) show the results for the dots and puzzle datasets, respectively, for a representative noise level (level 3 in previous experiments) and the Kendall tau distance. We can see that $\mathrm{OPT}^{d_{KT}}$ (solid red line) outperforms all other voting rules as long as $\widehat{t}$ is a reasonable overestimate of t^* ($r \in [1,2]$), but may or may not outperform them if $\widehat{t}$ is an underestimate of t^*. Again, other distance metrics generate similar results (see the appendix for details).

Comments on the empirical results. It is genuinely surprising that on real-world datasets, OPT^d (a rule designed to work well in the worst-case) provides a significantly superior average-case performance compared to most prominent voting rules by utilizing minimal additional information — an approximate upper bound on the average error in the input votes.

The inferior performance of methods based on probabilistic models of error is also thought provoking. After all, these models assume independent errors in the input votes, which is a plausible assumption in crowdsourcing settings. But such probabilistic models typically assume a specific structure on the distribution of the noise, e.g., the exponential distribution in Mallows' model [139], and it is almost impossible that noise in practice would follow this exact structure. In contrast, our approach only requires a loose upper bound on the average error in the input votes. In crowdsourcing settings where the noise is highly unpredictable, it can be argued that the principal may not be able to judge the exact distribution of errors, but may be able to provide an approximate bound on the average error.

9.6 Related Work

Our work is related to the extensive literature on error-correcting codes that use permutations (see, e.g., [17], and the references therein), but differs in one crucial aspect. In designing error-correcting codes, the focus is on two choices: i) the *codewords*, a subset of rankings which represent the "possible ground truths", and ii) the *code*, which converts every codeword into the message to be sent. These choices are optimized to achieve the best tradeoff between the number of errors corrected and the rate of the code (efficiency), while allowing unique identification of the ground truth. In contrast, our setting has fixed choices: i) every ranking is a possible ground truth, and ii) in coding theory terms, our setting constrains us to the repetition code. Both restrictions (inevitable in our setting) lead to significant inefficiencies, as well as the impossibility of unique identification of the ground truth (as illustrated in the introduction). Our research question is reminiscent of coding theory settings where a bound on adversarial noise is given, and a code is chosen with the bound on the noise as an input to maximize efficiency (see, e.g., [109]).

List decoding (see, e.g., [107]) relaxes classic error correction by guaranteeing that the number of possible messages does not exceed a small quota; then, the decoder simply lists all possible messages. The motivation is that one can simply scan the list and find the correct message, as all other messages on the list are likely to be gibberish. In the voting context, one cannot simply disregard some potential ground truths as nonsensical; we therefore select a ranking that is close to every possible ground truth.

Our model is also reminiscent of the distance rationalizability framework from the social choice literature [143]. In this framework, there is a fixed set of "consensus profiles" that admit an obvious output. Given a profile of votes, one finds the closest consensus profile (according to some metric), and returns the obvious output for that profile. Our model closely resembles the case where the consensus profiles are strongly unanimous, i.e., they consist of repetitions of a single ranking (which is also the ideal output). The key difference in our model is that instead of focusing solely on the closest ranking (strongly unanimous profile), we need to consider all rankings up to an average distance of t from the given profile — as they are all plausible ground truths — and return a single ranking that is at distance at most k from all such rankings.

A bit further afield, Procaccia et al. [176] study a probabilistic noisy voting setting, and quantify the *robustness* of voting rules to random errors. Their results focus on the probability that the outcome would change, under a random transposition of two adjacent alternatives in a single vote from a submitted profile, in the *worst-case* over profiles. Their work is different from ours in many ways, but perhaps most importantly, they are interested in how frequently common voting rules make mistakes, whereas we are interested in the guarantees of *optimal* voting rules that *avoid* mistakes.

9.7 Discussion

Uniformly accurate votes. Motivated by crowdsourcing settings, we considered the case where the average error in the input votes is guaranteed to be low. Instead, suppose we know that every vote in the input profile π is at distance at most t from the ground truth σ^*, i.e., $\max_{\sigma \in \pi} d(\sigma, \sigma^*) \leqslant t$. If t is small, this is a stronger assumption because it means that there are no outliers, which is implausible in crowdsourcing settings but plausible if the input votes are expert opinions. In this setting, it is immediate that any vote in the given profile is at distance at most $d^{\downarrow}(t)$ from the ground truth. Moreover, the proof of Theorem 9.4 goes through, so this bound is tight *in the worst case*; however, returning a ranking from the profile is *not optimal for every profile*.

Randomization. We did not consider randomized rules, which may return a distribution over rankings. If we take the error of a randomized rule to be the expected distance of the returned ranking from the ground truth, it is easy to obtain an upper bound of t. Again, the proof of Theorem 9.4 can be extended to yield an almost matching lower bound of $d^{\downarrow}(t)$. While randomized rules provide better guarantees, they are often impractical: low error is only guaranteed when rankings are repeatedly selected from the output distribution of the randomized rule on the same profile; however, most social choice settings see only a single outcome realized.[5]

Complexity. A potential drawback of the proposed approach is computational complexity. For example, consider the Kendall tau distance. When t is small enough, only the Kemeny ranking would be a possible ground truth, and $\text{OPT}^{d_{KT}}$ or any finite approximation thereof must return the Kemeny ranking, if it is unique. The $\mathcal{NP}$-hardness

[5]Exceptions include cases where randomization is used for circumventing impossibilities [34, 61, 168].

of computing the Kemeny ranking [20] therefore suggests that computing or approximating $\mathrm{OPT}^{d_{KT}}$ is $\mathcal{NP}$-hard.

One way to circumvent this computational obstacle is picking a ranking from the given profile, which provides a weaker bound of $3t$ instead of $2t$ on the distance from the unknown ground truth (see Theorem 9.2). However, in practice the optimal ranking can also be computed using various fixed-parameter tractable algorithms, integer programming solutions, and other heuristics, which are known to provide good performance for these types of computational problems (see, e.g., [23, 24]). More importantly, the crowdsourcing settings that motivate our work inherently restrict the number of alternatives to a relatively small constant: a human would find it difficult to effectively rank more than, say, 10 alternatives. With a constant number of alternatives, we can simply enumerate all possible rankings in polynomial time, making each and every computational problem considered in this chapter tractable. In fact, this is what we did in our experiments. Therefore, we do not view computational complexity as an insurmountable obstacle.

Appendix A

Omitted Proofs and Results for Chapter 2

A.1 The MNW Solution

The MNW solution is formally described in Algorithm 9. Computing an MNW solution consists of two stages: i) finding a largest set of players S to which one can simultaneously provide a positive utility, and ii) finding an allocation of the goods to players in S that maximizes their product of utilities. The implementation of the latter stage is described in detail in Section 2.5. For the former stage, S can be computed as follows.

Create a bipartite graph G with the players on one side and the goods on the other, and add an edge from player i to good g iff $v_i(g) > 0$.[1] For additive valuations — and more generally, for submodular valuations — it holds that if a player has a positive value for a bundle of goods, there must exist a good g in the bundle such that the player has a positive value for the singleton set $\{g\}$. Thus, at least for submodular valuations, to provide a positive utility to the maximum number of players it is sufficient to restrict our attention to allocations that assign at most one good to each player, i.e., represent a matching in the graph G. Our desired set S can now be computed as the set of players satisfied under a maximum cardinality matching in G. There are many popular polynomial time algorithms that one can use to find a maximum cardinality matching in a bipartite graph, e.g., the Hopcroft-Karp method.

Finally, we remark that while the set S can be computed in polynomial time for submodular (and thus for additive) valuations, this may be a computationally hard task for other classes of valuation functions.

A.2 Implementation on Spliddit

Section 2.5 outlines an implementation of the MNW solution when all the goods are indivisible. In contrast, our fair division website Spliddit allows an arbitrary mix of

[1] Recall that $v_i(g)$ is shorthand for $v_i(\{g\})$.

ALGORITHM 9: The MNW solution

Input: The set of players $\mathcal{N}$, the set of indivisible goods $\mathcal{M}$, and players' valuations $\{v_i\}_{i \in \mathcal{N}}$
Output: An MNW allocation $\vec{A}^{\text{MNW}}$

$S \in \underset{\substack{T \subseteq \mathcal{N} \text{ such that} \\ \exists \vec{A} \in \Pi_n(\mathcal{M}),\ \forall i \in T,\ v_i(A_i) > 0}}{\arg\max} |T|$; `// a largest set of players that can be simultaneously given a positive utility`

$\vec{A}^* \leftarrow \arg\max_{\vec{A} \in \Pi_{|S|}(\mathcal{M})} \prod_{i \in S} v_i(A_i)$; `// The MNW allocation to players in` S

$\vec{A}_i^{\text{MNW}} \leftarrow A_i^*, \forall i \in S$;

$\vec{A}_i^{\text{MNW}} \leftarrow \emptyset, \forall i \in \mathcal{N} \setminus S$; `// Players in` $\mathcal{N} \setminus S$ `do not receive goods`

divisible and indivisible goods, for which we designed an implementation that builds on the implementation of Section 2.5.

Splitting divisible goods:

As described in Section 2.7, one approach is to split each divisible good into k identical indivisible goods, and apply the MNW solution on the resulting set of indivisible goods. When k goes to infinity, this approach perfectly simulates the divisible goods, and gives the following relaxation of EF in addition to Pareto optimality (PO):

> For every pair of players i and j, there exists an indivisible good in player j's bundle such that player i does not envy player j after removing it from player j's bundle.

However, splitting each divisible good into infinitely many indivisible goods is computationally not feasible. In practice, it suffices to split each divisible good into 100 indivisible goods, which provides the following relaxation of EF in addition to PO:

> For every pair of players i and j, there exists either an indivisible good or 1% of a divisible good in player j's bundle such that player i does not envy player j after removing it from player j's bundle.

Final implementation:

Explicitly splitting each divisible good into 100 identical indivisible goods results in two computational challenges:

1. The number of goods, and, as a result, the number of decision variables in the resulting MILP increase significantly.
2. The number of constraints required to encode the piecewise-linear approximation of the logarithm function (in the form of segments or tangents on the log curve) is proportional to the number of possible utility levels that a player can achieve, which also increases from 1000 to 1000×100.

The former can be alleviated almost completely. Recall that the first step to computing the MNW solution is to find a largest set of players that can simultaneously derive a positive utility. This requires computing a maximum-cardinality matching, for which we use the MatlabBGL library.[2] Since the maximum-cardinality algorithm works on

[2] https://www.cs.purdue.edu/homes/dgleich/packages/matlab_bgl

sparse graphs and is extremely fast in practice, the increased number of goods is not an issue in this step.

The next step is to compute the MNW solution for the reduced set of players using the MILP of Figure 2.4. Here, the increased number of goods could affect the running time significantly. However, note that the indivisible goods created from a divisible good g are identical. Hence, we can retain the original decision variables $x_{i,g}$, but use them to denote the number of parts (out of 100) of good g that player i receives, rather than denoting whether player i receives good g entirely. In particular, for each divisible good g and each player i, we replace all the occurrences of $x_{i,g}$ in the MILP of Figure 2.4 with $x_{i,g}/100$, and replace $x_{i,g} \in \{0,1\}$ with $x_{i,g} \in \{0,1,\ldots,100\}$. The resulting MILP still has $n \cdot m$ integer (though, not binary) variables and n continuous variables, and we solve it using CPLEX.

Finally, for the latter challenge, note that although the number of possible utility levels that a player can achieve could, in the worst case, be 10^5, in practice it is significantly smaller. We use a preprocessing step to identify the possible utility levels for each player using a variant of the standard dynamic programming algorithm for the Knapsack problem, implemented efficiently in MATLAB through vectorization.

A.3 The Elusive Combination of EF1 and PO

In this section, we provide examples of several candidate solutions that fail to achieve EF1 and PO together for additive valuations — two properties that are fairly easy to achieve individually. This serves as a backdrop to our argument that it is compelling — even surprising — that the MNW solution achieves the two properties together (Theorem 2.1).

Example A.1 (Rounding any MNW allocation for divisible goods violates EF1)**.** The example we provide requires only 3 players but 31 goods. Let the set of players be $\mathcal{N} = \{1,2,3\}$. Suppose we have four types of goods: a single good of type a, and 10 goods each of types b, c, and d. Each player identically values all goods of the same type. Let the valuations of the players (specified only as a function of the type of the good) be as follows:

	Type a	Type b	Type c	Type d
Player 1	20	1	1.3	1.3
Player 2	15	0	1	1.3
Player 3	10	0	0	1

Using the KKT conditions, one can check that the unique MNW allocation when all goods are divisible is as follows. All the goods of type b, c, and d are allocated entirely to players 1, 2, and 3, respectively. The single good of type a is divided between the players such that players 1, 2, and 3 receive a 10/18, 7/18, and 1/18 fraction of the good, respectively.

Let us now find an allocation for *indivisible* goods by rounding this MNW allocation for divisible goods. Because the allocation for divisible goods does not divide goods of types b, c, and d, no rounding scheme can alter the allocation of these goods. However, we now show that subject to this constraint, allocating the single good of type a entirely to any single player violates EF1. Indeed, if we allocate the good to player 1 (resp. player 2), player 2 (resp. player 1) envies player 3 even after removing any single good from player 3's bundle. If we allocate the good to player 3, player 1 envies player 2 even after removing any single good from player 2's bundle.

This shows that in this example, no rounding scheme applied to the unique MNW allocation for divisible goods can produce an EF1 allocation of indivisible goods. Because Theorem 2.1 asserts that an MNW allocation of indivisible goods is guaranteed to be EF1 and PO, this is also a fascinating example in which no way of rounding the MNW allocation for divisible goods produces an MNW allocation for indivisible goods. In other words, an MNW allocation for indivisible goods inevitably gives at least one good to a player that receives a *zero* fraction of that good under the MNW solution for divisible goods.

In the economics literature, three popular notions of welfare — utilitarian, Nash, and egalitarian — are often arranged on a spectrum in which maximizing the utilitarian welfare is considered the most efficient, maximizing the egalitarian welfare is considered the fairest, and maximizing the Nash welfare is considered a good tradeoff between efficiency and fairness. While at first glance this interpretation may seem true in our setting as well — maximizing the Nash welfare does achieve both fairness (EF1) and efficiency (PO) — note that there is no "tradeoff" because, as the next example shows, maximizing either of the two other welfare notions does not guarantee EF1. From this axiomatic viewpoint, maximizing the Nash welfare in fact leads to a fairer outcome than maximizing either one of the other notions.

Example A.2 (Maximizing the utilitarian or the egalitarian welfare violates EF1)**.** The fact that maximizing the utilitarian welfare violates EF1 is very easy to see. Let the set of players be $\mathcal{N} = \{1, 2\}$, the set of goods be $\mathcal{M} = \{g_1, g_2, g_3\}$, and the additive valuations of the players be as follows:

	g_1	g_2	g_3
Player 1	1/2	1/2	0
Player 2	2/5	2/5	1/5

Note that the unique allocation that maximizes the utilitarian welfare allocates goods g_1 and g_2 to player 1, and good g_3 to player 2, causing player 2 to envy player 1 even after removal of any single good from player 1's bundle.

To show that maximizing the egalitarian welfare violates EF1, we use a slightly more involved example. Let the set of players be $\mathcal{N} = \{1, 2, 3\}$, the set of goods be $\mathcal{M} = \{g_1, g_2, g_3, g_4\}$, and the additive valuations of the players be as follows:

	g_1	g_2	g_3	g_4
Player 1	1	0	0	0
Player 2	2/3	0	1/6	1/6
Player 3	0	1/5	2/5	2/5

First, to achieve a positive egalitarian welfare we must allocate good 1 to player 1. Subject to this, the egalitarian welfare is uniquely maximized when good g_2 is allocated to player 3, and both goods g_3 and g_4 are allocated to player 2. However, this causes player 3 to envy player 2 even after removal of any single good from player 2's bundle.

Example A.3 (Maximizing the utilitarian/egalitarian welfare subject to EF1)**.** The following counterexample shows that maximizing the utilitarian welfare subject to EF1 violates PO. This example was discovered using computer simulations. Let the set of players be $\mathcal{N} = \{1, 2, 3, 4\}$, the set of goods be $\mathcal{M} = \{g_i\}_{i \in [10]}$, and the additive valuations of the players be as follows:

	g_1	g_2	g_3	g_4	g_5	g_6	g_7	g_8	g_9	g_{10}
1	0.0426	0.0004	0.1019	0.1503	0.0541	0.1782	0.1212	0.0259	0.1574	0.1681
2	0.0365	0.0004	0.2311	0.1479	0.0649	0.1150	0.1501	0.1894	0.0285	0.0362
3	0.1124	0.0972	0.0574	0.0956	0.1441	0.1461	0.0674	0.1272	0.0254	0.1273
4	0.0368	0.0582	0.0242	0.0784	0.1844	0.1260	0.1124	0.1121	0.1610	0.1064

It can be checked that maximizing the utilitarian welfare subject to the EF1 constraint results in the following allocation $\vec{A}$:

$$A_1 = \{g_6, g_7, g_{10}\}, A_2 = \{g_3, g_4, g_8\}, A_3 = \{g_1, g_2\}, \text{ and } A_4 = \{g_5, g_9\}.$$

However, this allocation is not Pareto optimal. An alternative allocation in which players 1 and 2 exchange goods g_7 and g_4 improves the utility to both players 1 and 2 while keeping the utility to both players 3 and 4 unaltered. This alternative allocation is not selected in the first place because it violates EF1 (player 3 now envies player 1 even after the removal of any single good from player 1's bundle).

It is easy to see why maximizing the egalitarian welfare subject to EF1 violates PO. Suppose the set of players is $\mathcal{N} = \{1, 2, 3\}$, and the set of goods is $\mathcal{M} = \{g_1, g_2, g_3\}$. Let the valuations of the players be as follows:

	g_1	g_2	g_3
Player 1	2/3	1/3	0
Player 2	1/3	2/3	0
Player 3	1/3	1/3	1/3

Clearly the optimal egalitarian welfare is 1/3 in this example. An EF1 allocation $\vec{A}$ that achieves this optimal welfare is given by $A_1 = \{g_2\}$, $A_2 = \{g_1\}$, and $A_3 = \{g_3\}$. However, this is clearly not PO: if players 1 and 2 exchange their bundles, they can both be better off without reducing the utility to player 3. Hence, maximizing the egalitarian welfare (at least naïvely) subject to EF1 is not PO.

While EF1 and PO are both mild properties by themselves, their combination is surprisingly elusive, which provides a justification for the MNW solution. However, EF1 and PO by themselves are not sufficient to always guarantee a desirable outcome. The following observations illustrate why.

Example A.4. Imagine we have a set of two players $\mathcal{N} = \{1, 2\}$, and a set of two goods $\mathcal{M} = \{g_1, g_2\}$. Suppose player 1 values both goods equally, and player 2 only values good g_2.

In this case, the only intuitively fair outcome (which is also the outcome that the MNW solution selects) assigns good g_1 to player 1, and good g_2 to player 2. However, note that assigning both goods to player 1 also satisfies EF1 and PO, but is clearly undesirable.

More formally, we can argue that, while the MNW solution provides $\pi_n = 1/\Theta(\sqrt{n})$-approximation of the MMS guarantee, simply restricting the allocation to be EF1 and PO gives a worse $1/n$-approximation of MMS.

Theorem A.1. *Every allocation that is envy free up to one good (EF1) and Pareto optimal (PO) is $1/n$-maximin share (MMS) for additive valuations over indivisible goods. Further, the factor $1/n$ is tight, i.e., for every $n \in \mathbb{N}$ and $\varepsilon > 0$, there exists an instance with n players having additive valuations and an allocation satisfying EF1 and PO that is not $(1/n + \varepsilon)$-MMS.*

Proof. We first prove that every allocation satisfying EF1 and PO is $1/n$-MMS, and later prove tightness of the approximation ratio (upper bound).

Proof of the lower bound: Let $\vec{A}$ be an allocation satisfying EF1 and PO. Fix a player $i \in \mathcal{N}$. Because $\vec{A}$ is EF1, for every player $j \in \mathcal{N} \setminus \{i\}$ there exists a good $g_{ij} \in A_j$ such that

$$v_i(A_i) \geqslant v_i(A_j) - v_i(g_{ij}). \tag{A.1}$$

Let $T_i = \sum_{g \in \mathcal{M}} v_i(g)$, and let $E_i = \sum_{j \in \mathcal{N} \setminus \{i\}} v_i(g_{ij})$. Then, summing Equation (A.1) over all $j \in \mathcal{N} \setminus \{i\}$, we get

$$(n-1) \cdot v_i(A_i) \geqslant \textstyle\sum_{j \in \mathcal{N} \setminus \{i\}} v_i(A_j) - E_i \Rightarrow n \cdot v_i(A_i) \geqslant T_i - E_i. \tag{A.2}$$

On the other hand, consider any partition of the set of goods $\mathcal{M}$ into n bundles. Due to the pigeonhole principle, there must exist a bundle that does not contain good g_{ij} for any $j \in \mathcal{N} \setminus \{i\}$. Since the value of this bundle according to player i can be at most $T_i - E_i$, the MMS guarantee of player i is also at most $T_i - E_i$. Equation (A.2) now implies that under $\vec{A}$, each player i receives at least $1/n$ of her MMS guarantee, i.e., $\vec{A}$ is $1/n$-MMS.

Proof of the upper bound: We now show that for every $n \in \mathbb{N}$ and $\varepsilon > 0$, there exists an instance with n players for which some allocation satisfying EF1 and PO is not $(1/n + \varepsilon)$-MMS.

Construct an instance with n players and $2n-1$ goods. Let there be $n-1$ "high" goods that each player values at n, and n "low" goods that each player values at 1. The MMS guarantee of each player is n: the player can put each "high" good in its own bundle, and all "low" goods in a single bundle.

However, one can check that giving $n-1$ of the players a high and a low good each, and giving the remaining player the remaining single low good also satisfies EF1 and PO, but gives the last player exactly $1/n$ of her MMS guarantee. ■

A.4 General Valuations

In this section, we provide the definitions of the families of valuation functions mentioned in Section 2.3.1, and provide the missing proofs and examples. Let us begin by formally defining subadditive, superadditive, submodular, and supermodular valuations.

Definition A.1 (Subadditive and Superadditive Valuations)**.** A valuation function $v : 2^{\mathcal{M}} \to \mathbb{R}_{\geqslant 0}$ is called subadditive (resp. superadditive) if for every pair of disjoint sets $S, T \subseteq \mathcal{M}$, we have $v(S \cup T) \leqslant v(S) + v(T)$ (resp. $v(S \cup T) \geqslant v(S) + v(T)$).

Definition A.2 (Submodular and Supermodular Valuations)**.** A valuation function $v : 2^{\mathcal{M}} \to \mathbb{R}_{\geqslant 0}$ is called submodular (resp. supermodular) if for every pair $S, T \subseteq \mathcal{M}$, we have $v(S \cup T) \leqslant v(S) + v(T) - v(S \cap T)$ (resp. $v(S \cup T) \geqslant v(S) + v(T) - v(S \cap T)$).

It is clear that submodular (resp. supermodular) valuations are a special case of subadditive (resp. superadditive) valuations. We now provide a proof of Theorem 2.2, which asserts that for supermodular (and thus superadditive) valuations and subadditive valuations, EF1 and PO are incompatible.

of Theorem 2.2. Let the set of players be $\mathcal{N} = \{1, 2\}$, and the set of goods be $\mathcal{M} = \{a, b, c, d\}$. We use a common valuation for both players. Figures A.2 and A.1 define the supermodular (thus superadditive) valuation v^{sup} and the subadditive valuation v^{sub}, respectively, through their value for a set $S \subseteq \mathcal{M}$.

$$v^{\text{sub}}(S) = \begin{cases} 10 & \text{if } |S| = 4, \\ 7 & \text{if } |S| = 3, \\ 6 & \text{if } |S| = 2 \text{ and } a \notin S, \\ 4 & \text{if } |S| = 2 \text{ and } a \in S, \\ 4 & \text{if } S = \{a\}, \\ 3 & \text{if } S = \{b\}, \{c\}, \text{or } \{d\}, \\ 0 & \text{if } S = \emptyset. \end{cases}$$

Figure A.1: Subadditive valuation

$$v^{\text{sup}}(S) = \begin{cases} 4 & \text{if } |S| = 4, \\ 3 & \text{if } |S| = 3, \\ 2 & \text{if } |S| = 2 \text{ and } a \notin S, \\ 1 & \text{if } |S| = 2 \text{ and } a \in S, \\ 1 & \text{if } S = \{a\}, \\ 0 & \text{if } S = \{b\}, \{c\}, \{d\}, \text{or } \emptyset. \end{cases}$$

Figure A.2: Supermodular (thus superadditive) valuation

In each case, under a PO allocation, player 1 receives one of the following sets of goods: $\emptyset$, $\{a\}$, $\{b, c, d\}$, and $\mathcal{M}$; and player 2 receives the set of remaining goods. It is easy to check that these are the only four PO allocation. Note that EF1 is violated in the first two allocations due to player 1 envying player 2 (and in last two allocations due to player 2 envying player 1) even after removal of any single good from the envied player's bundle. ■

We now focus on the interesting case of submodular valuations Submodular valuations are characteristic of substitute goods, and are alternatively defined via diminishing marginal utility. Examples of submodular valuations include unit demand valuations, strong valuations with no complementarities, and gross substitutes.

As mentioned in Section 2.3.1, we were unable to settle the question of the compatibility of EF1 and PO for submodular valuations. We know that an MNW allocation does not guarantee EF1 and PO for submodular valuations, but we can show that it guarantees MEF1 (a relaxation of EF1 that coincides with EF1 for additive valuations) and PO.

Example A.5 (MNW is not EF1 and PO for submodular valuations)**.** Let the set of players be $\mathcal{N} = \{1, 2\}$, and the set of goods be $\mathcal{M} = \{a, b, c, d\}$. The submodular valuations v_1 and v_2 of players 1 and 2, respectively, are as follows:
Player 1*:* First, let us define the value of the player for individual goods.

$$v_1(a) = 1, v_1(b) = 1, v_1(c) = 0, \text{ and } v_1(d) = 0.$$

For $S \subseteq \mathcal{M}$ with $|S| \geqslant 2$, define $v_1(S)$ to be the sum of the values of the *two goods* in S that are the most valuable to the player. It is easy to check that this is a submodular valuation.
Player 2*:* Let the value of the player for individual goods be as follows.

$$v_2(a) = 2.5, v_2(b) = 2.5, v_2(c) = 1, \text{ and } v_2(d) = 1.$$

Once again, for $S \subseteq \mathcal{M}$ with $|S| \geqslant 2$, define $v_2(S)$ to be the sum of the values of the *two goods* in S that are the most valuable to the player. Similarly to v_1, v_2 is also a submodular valuation.

Note that an MNW allocation must allocate at least one of the two goods a or b to player 1 to achieve positive Nash welfare. If player 1 receives only one of these two goods, the Nash welfare can be at most $1 \cdot 3.5 = 3.5$. In contrast, the allocation that gives goods a and b to player 1, and goods c and d to player 2, achieves Nash welfare of $2 \cdot 2 = 4$. Hence, this allocation is the unique MNW allocation. However, player 2 then envies player 1 even after removal of any single good from player 1's bundle.

We end this section with a proof of Theorem 2.3, which asserts that every MNW allocation is MEF1 and PO for submodular valuations.

of Theorem 2.3. Let $\vec{A}$ be an MNW allocation. First, let us prove the result for the case of $\mathtt{NW}(\vec{A}) > 0$. In this case, the Pareto optimality of $\vec{A}$ is obvious due to the fact that $\vec{A}$ maximizes the Nash welfare. Suppose, for contradiction, that $\vec{A}$ is not MEF1. Then, there exist players $i, j \in N$ such that

$$\forall g \in A_j, v_i(A_i \cup A_j \setminus \{g\}) - v_i(A_i) > v_i(A_i). \tag{A.3}$$

Next, for every $r \in A_j$, let us define

$$\delta_i(g) = v_i(A_i \cup \{g\}) - v_i(A_i), \text{ and } \delta_j(g) = v_j(A_j) - v_j(A_j \setminus \{g\}).$$

Note that $\delta_i(g)$ and $\delta_j(g)$ are generalizations of $v_i(g)$ and $v_j(g)$ from additive valuations to submodular valuations. Also, observe that they are defined a bit differently for i and j.

We now derive two key results.

Lemma A.1. *For every $g^* \in A_j$, we have $\sum_{g \in A_j} \delta_i(g) > v_i(A_i \cup \{g^*\})$.*

Proof. Fix $g^* \in A_j$. Let us enumerate the elements of A_j as $g_1, \ldots, g_k$ where $k = |A_j|$ and $g_k = g^*$. Also, for $t \in [k]$ define $A_j^t = \{g_1, \ldots, g_t\}$, and $A_j^0 = \emptyset$. Then,

$$\sum_{g \in A_j \setminus \{g^*\}} \delta_i(g) = \sum_{t=1}^{k-1} v_i(A_i \cup \{g_t\}) - v_i(A_i) \geqslant \sum_{t=1}^{k-1} v_i(A_i \cup A_j^t) - v_i(A_i \cup A_j^{t-1})$$
$$= v_i(A_i \cup A_j \setminus \{g^*\}) - v_i(A_i) > v_i(A_i),$$

where the second transition holds due to submodularity of v_i and the final transition follows from Equation (A.3). Adding $\delta_i(g^*) = v_i(A_i \cup \{g^*\}) - v_i(A_i)$ on both sides yields the desired result. ■ (Proof of Lemma A.1)

Lemma A.2. *We have $\sum_{g \in A_j} \delta_j(g) \leqslant v_j(A_j)$.*

Proof. Once again, let $A_j = \{g_1, \ldots, g_k\}$, where $k = |A_j|$, $A_j^t = \{g_1, \ldots, g_t\}$ for $t \in [k]$, and $A_j^0 = \emptyset$. Then,

$$\sum_{g \in A_j} \delta_j(g) = \sum_{t=1}^{k} v_j(A_j) - v_j(A_j \setminus \{g_t\}) \leqslant \sum_{t=1}^{k} v_j(A_j^t) - v_j(A_j^{t-1}) = v_j(A_j),$$

where the inequality follows from the submodularity of v_j. ■ (Proof of Lemma A.2)

From Lemma A.1, it is clear that $\sum_{g \in A_j} \delta_i(g) > 0$. Thus, there exists $g \in A_j$ such that $\delta_i(g) > 0$. Fix $g^* = \arg\min_{g \in A_j : \delta_i(g) > 0} \delta_j(g)/\delta_i(g)$. We now take the ratio of the inequality in Lemma A.2 to the inequality in Lemma A.1 applied to our chosen g^*. This is well-defined because we already showed $\sum_{g \in A_j} \delta_i(g) > 0$, and we also have $v_i(A_i \cup \{g^*\}) \geqslant v_i(A_i) > 0$.

$$\frac{v_j(A_j)}{v_i(A_i \cup \{g^*\})} \geqslant \frac{\sum_{g \in A_j} \delta_j(g)}{\sum_{g \in A_j} \delta_i(g)} \geqslant \frac{\delta_j(g^*)}{\delta_i(g^*)} = \frac{v_j(A_j) - v_j(A_j \setminus \{g^*\})}{v_i(A_i \cup \{g^*\}) - v_i(A_i)},$$

where the second transition holds due to our choice of g^*. Upon rearranging the terms, we get

$$v_i(A_i \cup \{g^*\}) \cdot v_j(A_j \setminus \{g^*\}) > v_i(A_i) \cdot v_j(A_j),$$

which is a contradiction because it implies that shifting g^* from player j to player i would increase the Nash welfare, which is in direct violation of the optimality of the Nash welfare under the MNW allocation $\vec{A}$.

Let us now handle the case of $\texttt{NW}(\vec{A}) = 0$. Let S denote the set of players that receive positive utility under $\vec{A}$. The proof of Pareto optimality of $\vec{A}$ for submodular valuations is identical to the proof of Pareto optimality of an MNW allocation for additive valuations, which does not use additivity of the valuations. We now show that $\vec{A}$ is

MEF1. Note that MEF1 holds among players in S due to the proof of the previous case, and holds trivially among players in $\mathcal{N} \setminus S$. Hence, the only case we need to address is when a player $i \in \mathcal{N} \setminus S$ (with $A_i = \emptyset$) marginally envies player $j \in S$ (with $v_j(A_j) > 0$) up to one good. Then, by the definition of MEF1, we have

$$\forall g \in A_j, v_i(A_j \setminus \{g\}) > 0. \tag{A.4}$$

Submodularity of v_j implies that $\sum_{g \in A_j} v_j(\{g_j\}) \geqslant v_j(A_j) > 0$. Hence, there exists a good $\widehat{g} \in A_j$ such that $v_j(\{\widehat{g}\}) > 0$. Applying Equation (A.4) to $\widehat{g}$, we get $v_i(A_j \setminus \{\widehat{g}\}) > 0$. But then moving all goods in A_j except $\widehat{g}$ from player j to player i gives positive utility to player i while still giving positive utility to player j, which violates the fact that $\vec{A}$ provides positive utility to the maximum number of players. Hence, $\vec{A}$ must be MEF1. ■ (Proof of Theorem 2.3)

A.5 Pairwise Maximin Share Guarantee

In this section, we prove several results about our novel fairness concept — the pairwise maximin share guarantee.

Theorem A.2. *The pairwise maximin share guarantee is implied by envy-freeness (EF), and implies 1/2-maximin share guarantee, envy freeness up to the least valued good (EFX), and as a direct consequence, envy-freeness up to one good (EF1).*

Proof. Let $\vec{A}$ be an EF allocation, i.e., $v_i(A_i) \geqslant v_i(A_j)$ for all pairs of players $i, j \in \mathcal{N}$. Let $\texttt{PMMS}_i$ denote the pairwise MMS guarantee of player i:

$$\texttt{PMMS}_i = \max_{j \in \mathcal{N} \setminus \{i\}} \max_{\vec{B} \in \Pi_2(A_i \cup A_j)} \min\{v_i(B_1), v_i(B_2)\}.$$

Then, we have

$$\texttt{PMMS}_i \leqslant \max_{j \in \mathcal{N} \setminus \{i\}} \frac{v_i(A_i), v_i(A_j)}{2} \leqslant v_i(A_i),$$

where the first transition holds because its right hand side is the pairwise MMS guarantee that player i would have if all goods were divisible, which is an upper bound on $\texttt{PMMS}_i$ because divisible goods offer the player a greater flexibility in partitioning the goods. The second transition follows directly from the envy-freeness of $\vec{A}$.

Next, let $\vec{A}$ be a pairwise MMS allocation. It is easy to show that $\vec{A}$ must also be EFX: if player i envies player j after the removal of player i's least positively valued good g^* from A_j, then it follows that player i's pairwise MMS guarantee is at least $v_i(A_i \cup \{g^*\}) > v_i(A_i)$ due to the partition $(A_i \cup \{g^*\}, A_j \setminus \{g^*\})$. However, this implies that $\vec{A}$ is not pairwise MMS, which is a contradiction. Hence, $\vec{A}$ is also EFX. It is trivial to check that EFX implies EF1 by definition; hence, $\vec{A}$ is also EF1.

Finally, we show that a pairwise MMS allocation $\vec{A}$ is also 1/2-MMS. Consider players i and j. There are only two possible cases: (i) A_j has at most one good that player i

values positively, i.e., $|A_j \cap \{g \in \mathcal{M} \mid v_i(g) > 0\}| \leqslant 1$, or (ii) $v_i(A_j) \leqslant 2 \cdot v_i(A_i)$. Indeed, if A_j has at least two goods that player i values positively, and $v_i(A_j) > 2 \cdot v_i(A_i)$, then consider the good g^* that is the least valuable among player i's positively valued goods in A_j. In that case, player i could partition $A_i \cup A_j$ into $(A_i \cup \{g^*\}, A_j \setminus \{g^*\})$ and ensure that her pairwise MMS value is strictly more than $v_i(A_i)$, which is a contradiction because $\vec{A}$ is pairwise MMS.

Now, if no player in $\mathcal{N} \setminus \{i\}$ falls into case (ii), then it is easy to see that the MMS guarantee of player i is at most $v_i(A_i)$. If a non-empty subset $S \subseteq \mathcal{N} \setminus \{i\}$ of players fall into case (ii), then we can bound the MMS guarantee of player i from above by assuming that all goods allocated to players $S \cup \{i\}$ are divisible. However, this still gives an MMS guarantee of at most $2 \cdot v_i(A_i)$, because each player in $j \in S \cup \{i\}$ satisfies $v_i(A_j) \leqslant 2 \cdot v_i(A_i)$. Thus, the MMS guarantee of player i is at most $2 \cdot v_i(A_i)$, which implies that $\vec{A}$ is 1/2-MMS. ■

Finally, we give a proof of Corollary 2.1 that uses the MMS approximation guarantee of the MNW solution (Theorem 2.4) to prove a pairwise MMS approximation guarantee.

of Corollary 2.1. An MNW allocation $\vec{A}$ has the following interesting property: Take the goods allocated to players i and j, i.e., $\mathcal{M}' = A_i \cup A_j$, and take the set of players $\mathcal{N}' = \{i, j\}$. Then the allocation given by A_i and A_j is also an MNW allocation for the reduced instance of allocating the set of goods $\mathcal{M}'$ to the set of players $\mathcal{N}'$. This fact is easy to see when either $v_i(A_i) > 0$ and $v_j(A_j) > 0$ (otherwise we could achieve higher Nash welfare), or $v_i(A_i) = v_j(A_j) = 0$. When $v_i(A_i) = 0$ but $v_j(A_j) > 0$ (without loss of generality), every allocation of $\mathcal{M}'$ to players $\{i, j\}$ must provide zero utility to at least one player, otherwise this part of the allocation could be used in the original instance to increase the number of players that receive positive utility, contradicting the fact that an MNW allocation provides positive utility to the maximum number of players. Hence, the allocation in the reduced instance that provides all the goods in $\mathcal{M}'$ to player j (which is exactly allocation $\vec{A}$ restricted to the reduced instance) is indeed an MNW allocation, and is π_2-MMS in the reduced instance (Theorem 2.4).

We therefore conclude that the MNW allocation $\vec{A}$ is Φ-pairwise MMS in the original instance as $\pi_2 = \Phi$. To establish tightness of the factor Φ, for a given $n \in \mathbb{N}$ and $\varepsilon > 0$, we simply use the example from the proof of the upper bound in Theorem 2.4 after replacing π_n by $\pi_2 = \Phi$ in the valuations of the players. In the new example, now the pairwise MMS approximation ratio (instead of the MMS approximation ratio in the original example) can be driven below $\pi_2 + \varepsilon$ for a value of ε' less than $\min(\pi_2, \varepsilon/(\pi_2 \cdot (\pi_2 + \varepsilon)))$, which is a bound obtained by substituting $n = 2$ in the upper bound on ε' from the proof of Theorem 2.4. ■

A.6 A Spectrum of Fair and Efficient Solutions

In this work we focused on the MNW solution that maximizes the Nash welfare, i.e., selects an allocation $\vec{A}$ that maximizes $\prod_{i \in \mathcal{N}} v_i(A_i)$. For simplicity, let us assume $\mathcal{N} = [n]$.

Another popular solution concept for fair allocation, originally used in the networking literature, is *proportional fairness* [118].

Definition A.3 (Proportional Fairness)**.** An allocation $\vec{A}$ is said to satisfy proportional fairness if for any alternative allocation $\vec{A}'$, it holds that

$$\sum_{i=1}^{n} \frac{v_i(A'_i) - v_i(A_i)}{v_i(A_i)} \leqslant 0.$$

In words, an allocation is proportionally fair if the total percentage change in the players' utilities is non-positive when switching to any alternative allocation. This, in some sense, indicates that $\vec{A}$ is socially preferred over any alternative allocation.

Note that the proportional fairness requirement can equivalently be written as

$$\sum_{i=1}^{n} \frac{v_i(A'_i)}{v_i(A_i)} \leqslant n \Leftrightarrow \frac{n}{\sum_{i=1}^{n} \left(\frac{v_i(A_i)}{v_i(A'_i)}\right)^{-1}} \geqslant 1.$$

That is, proportional fairness requires that the *Harmonic mean* of the set of quantities $\{v_i(A_i)/v_i(A'_i)\}_{i \in [n]}$ be at least 1. Interestingly, the requirement for an allocation $\vec{A}$ to be an MNW allocation can be formulated in a similar manner, by requiring that the Nash welfare should not increase when switching to any alternative allocation $\vec{A}'$.

$$\prod_{i=1}^{n} v_i(A_i) \geqslant \prod_{i=1}^{n} v_i(A'_i) \Leftrightarrow \sqrt[n]{\prod_{i=1}^{n} \frac{v_i(A_i)}{v_i(A'_i)}} \geqslant 1.$$

That is, the MNW solution requires that the *geometric mean* of the same set of quantities $\{v_i(A_i)/v_i(A'_i)\}_{i \in [n]}$ be at least 1.

This inspired us to define a spectrum of properties for the allocation of indivisible goods which require that the p-th power mean of the same set of quantities be at least 1. Recall that the p-th power mean of a set of non-negative numbers $\{x_i\}_{i \in [n]}$ is defined as $\left(\frac{1}{n} \sum_{i=1}^{n} x_i^p\right)^{1/p}$. The Harmonic mean corresponds to $p = -1$, and the geometric mean corresponds to $p = 0$. To be symmetric, we define the property up to $p = 1$.

Definition A.4. For $p \in [-1, 1]$, we say that an allocation $\vec{A} \in \Pi_n(\mathcal{M})$ satisfies $\Gamma(p)$ if for every other allocation $\vec{A}' \in \Pi_n(\mathcal{M})$, we have

$$\left(\frac{1}{n} \sum_{i=1}^{n} \left(\frac{v_i(A_i)}{v_i(A'_i)}\right)^p\right)^{1/p} \geqslant 1.$$

We now make several interesting observations about the properties of the allocations that satisfy $\Gamma(p)$. First, the power-mean inequality states that for $p < p'$, the p-th power mean is no more than the p'-th power mean. This directly yields that $\Gamma(p)$ implies $\Gamma(p')$.

Theorem A.3 (Decreasing Power). *For $p, p' \in [-1, 1]$ with $p' > p$, every allocation satisfying $\Gamma(p)$ also satisfies $\Gamma(p')$.*

We know that there always exists an allocation (any MNW allocation) that satisfies $\Gamma(0)$, and therefore $\Gamma(p)$ for all $p > 0$. In contrast, there exist instances in which no allocation satisfies $\Gamma(p)$ for any $p < 0$; simply consider a single good and two players having value 1 for the good.

Theorem A.4 (Existence). *For $p \in [-1, 1]$, an allocation satisfying $\Gamma(p)$ always exists if and only if $p \geqslant 0$.*

We now show that $\Gamma(p)$ implies important efficiency and fairness properties.

Theorem A.5 (Efficiency). *For $p \in [-1, 1]$, every allocation satisfying $\Gamma(p)$ is Pareto optimal (PO).*

Proof. Indeed, assume that an allocation $\vec{A} \in \Pi_n(\mathcal{M})$ satisfying $\Gamma(p)$ is not PO. Let $\vec{A}' \in \Pi_n(\mathcal{M})$ be another allocation that satisfies $v_i(A'_i) \geqslant v_i(A_i)$ for all $i \in \mathcal{N}$, and $v_{i^*}(A'_{i^*}) > v_{i^*}(A_{i^*})$ for some $i^* \in \mathcal{N}$. Then, we would have

$$\sum_{i=1}^{n} \left(\frac{v_i(A_i)}{v_i(A'_i)} \right)^p \quad < \quad n,$$

which is a contradiction because $\vec{A}$ satisfies $\Gamma(p)$. ■

We studied a relaxation of the ideal fairness notion envy-freeness (EF), called envy-freeness up to one good (EF1). One can similarly define *envy-freeness up to k goods*: An allocation $\vec{A}$ is envy-free up to k goods if

$$\forall i, j \in \mathcal{N}, \exists S \subseteq A_j \text{ with } |S| \leqslant k, v_i(A_i) \geqslant v_i(A_j \setminus S).$$

Theorem A.6 (Fairness). *For $p \in [-1, 1]$, every allocation satisfying $\Gamma(p)$ is envy free up to $1 + \lceil p \rceil$ goods, where $\lceil \cdot \rceil$ is the ceiling function.*

Proof. Due to Theorem A.3, we only need to prove this theorem for $p \in \{-1, 0, 1\}$. For $p = 0$, we already showed that every MNW allocation is EF1 (Theorem 2.1).

Let us now consider $p = -1$. Let allocation $\vec{A} \in \Pi_n(\mathcal{M})$ satisfy $\Gamma(-1)$. Consider a pair of players j, j'. For every good $t \in A_{j'}$, we apply the inequality in the definition of $\Gamma(-1)$ using the allocation $\vec{A}$ and the allocation $\vec{A}'$ obtained by moving good t from player j' to player j. We have

$$n \quad \geqslant \quad \sum_{i=1}^{n} \frac{v_i(A'_i)}{v_i(A_i)} = \left(\sum_{i \neq j, j'} \frac{v_i(A_i)}{v_i(A_i)} \right) + \frac{v_j(A_j \cup \{t\})}{v_j(A_j)} + \frac{v_{j'}(A_{j'} \setminus \{t\})}{v_{j'}(A_{j'})} = n + \frac{v_j(t)}{v_j(A_j)} - \frac{v_{j'}(t)}{v_{j'}(A_{j'})}$$

which implies that $v_j(t) \leqslant v_{j'}(t) \cdot v_j(A_j) / v_{j'}(A_{j'})$. Summing over all $t \in A_{j'}$, we get

$$v_j(A_{j'}) \quad = \quad \sum_{t \in A_{j'}} v_j(t) \leqslant \sum_{t \in A_{j'}} v_{j'}(t) \frac{v_j(A_j)}{v_{j'}(A_{j'})} = v_j(A_j),$$

i.e., player j is not envious for player j'.

Let us now consider $p = 1$. Consider an allocation $\vec{A} \in \Pi_n(\mathcal{M})$ that satisfies $\Gamma(1)$. Consider players i and j. We want to show that player i would not envy j if we are allowed to remove (up to) two goods from player j's bundle. If $|A_j| \leqslant 2$, we are done. Assume $|A_j| \geqslant 3$. We now show that there are goods $t_1, t_2 \in A_j$ such that $v_i(A_i) \geqslant v_i(A_j \setminus \{t_1, t_2\}$.

Consider a good $t \in A_j$, and define the allocation $\vec{A}'$ obtained by moving good t from player j to player i in $\vec{A}$. By the definition of $\Gamma(1)$, we get

$$\frac{v_i(A_i)}{v_i(A_i) + v_i(k)} + \frac{v_j(A_j)}{v_j(A_j) - v_j(t)} \quad \geqslant \quad 2.$$

Setting $x_t = \frac{v_i(t)}{v_i(A_i)}$ and $y_t = \frac{v_j(t)}{v_j(A_j)}$, this inequality can be expressed as

$$\frac{1}{1 + x_t} + \frac{1}{1 - y_t} \quad \geqslant \quad 2$$

which implies that

$$x_t \quad \leqslant \quad \frac{y_t}{1 - 2y_t} \tag{A.5}$$

whenever $y_t \leqslant 1/3$. Now let t_1 and t_2 be the goods in A_j for which player j has the highest and second highest value, respectively. Hence, for every good $t \in A_j \setminus \{t_1, t_2\}$, we have $y_t \leqslant 1/3$. Using Equation (A.5), we obtain that the value of player i for the goods in $A_j \setminus \{t_1, t_2\}$ is

$$\begin{aligned}
v_i(A_j \setminus \{t_1, t_2\}) \quad &= \quad v_i(A_i) \sum_{t \in A_j \setminus \{t_1, t_2\}} \frac{v_i(t)}{v_i(A_i)} = v_i(A_i) \sum_{t \in A_j \setminus \{t_1, t_2\}} x_t \\
&\leqslant \quad v_i(A_i) \sum_{t \in A_j \setminus \{t_1, t_2\}} \frac{y_t}{1 - 2y_t} \leqslant v_i(A_i) \sum_{t \in A_j \setminus \{t_1, t_2\}} \frac{y_t}{1 - 2y_{t_2}} \\
&\leqslant \quad v_i(A_i),
\end{aligned}$$

where the second inequality holds because $y_t \leqslant y_{t_2}$ for $t \in A_j \setminus \{t_1, t_2\}$, and the final transition follows due to the definitions of y, t_1, and t_2. We thus have that $\vec{A}$ is EF2. ■

From Theorems A.4 and A.3, we know that the MNW solution, which satisfies $\Gamma(0)$, is the optimal solution on the spectrum that is guaranteed to exist as its fairness guarantee (EF1) is strictly better than the fairness guarantee provided by $\Gamma(p)$ for any $p > 0$. One may question whether the weaker $\Gamma(p)$ with $p > 0$ has a computational advantage over the MNW solution, which we know is $\mathcal{NP}$-hard [155]. Interestingly, a polynomial-time Turing reduction from the popular $\mathcal{NP}$-hard PARTITION problem shows that computing an allocation satisfying $\Gamma(p)$ is $\mathcal{NP}$-hard for $p \in (0, 1]$. Note that it is the search problem (of actually finding the allocation) that is $\mathcal{NP}$-hard rather than the decision problem of determining the existence of such an allocation (which is a trivial problem as such an allocation always exists).

Theorem A.7 (Computational Hardness)**.** *For $p \in [0,1]$, computing an allocation that satisfies property $\Gamma(p)$ is $\mathcal{NP}$-hard.*

Proof. Due to Theorem A.3, we only need to show the hardness for $p = 1$. We show that a polynomial-time algorithm to compute an allocation that is $\Gamma(1)$ can be used to decide the PARTITION problem in polynomial time. The input in an instance of the PARTITION problem is a set of m positive integers $S = \{x_1, \ldots, x_m\}$, and our goal is to decide whether there exists a *perfect partition* of S, i.e., a partition of S into two exclusive and exhaustive subsets whose sum of elements is equal. Let $T = \sum_{i \in [m]} x_i$. We say that a partition of S is a *minimum-difference partition* if the difference between the sums of the two subsets is the least possible among all partitions of S into two subsets.

Let us first construct a new set of $m' = m + 2$ positive integers $S' = \{x'_i\}_{i \in [m']}$ where $x'_i = 5x_i$ for $i \in [m]$, $x_{m+1} = 1$, and $x_{m+2} = 2$. Let $T' = \sum_{i \in [m']} x'_i = 5T + 3$. Note that S' does not have a perfect partition. Further, it has a minimum-difference partition with difference 1 if and only if S has a perfect partition. Note that a partition of S' with difference 1 can only be created by taking a perfect partition of S, replacing the elements of S by the corresponding elements of S', and then adding x'_{m+1} and x'_{m+2} in different subsets.

Next, we construct an instance of our fair allocation problem as follows. We have two players with the identical valuation v over the set of goods $\mathcal{M} = [m']$ under which $v(i) = x'_i$ for each good $i \in \mathcal{M}$. We can interpret an allocation $\vec{A}$ of this instance as a partition of S', in which each subset is formed by taking the elements of S' corresponding to the goods in a player's bundle. Thus, the sums of the two subsets in the partition are exactly $v(A_1)$ and $v(A_2)$, and $v(A_1) + v(A_2) = T'$.

We now show that every allocation satisfying $\Gamma(1)$ produces a minimum-difference partition of S'. To see this, consider an allocation $\vec{A}$ satisfying $\Gamma(1)$, and without loss of generality, assume $v(A_1) - v(A_2) = \delta > 0$. Thus, $v(A_1) = (T' + \delta)/2$ and $v(A_2) = (T' - \delta)/2$. Now, suppose for contradiction that there exists another allocation $\vec{A}'$ under which $|v(A'_1) - v(A'_2)| = \varepsilon < \delta$. Because S' does not admit a perfect partition, we have $\varepsilon > 0$. Without loss of generality, let $v(A'_1) - v(A'_2) = \varepsilon$ (otherwise we can switch the bundles of the two players). Hence, $v(A'_1) = (T + \varepsilon)/2$ and $v(A'_2) = (T - \varepsilon)/2$. However, in this case

$$\begin{aligned}
&\frac{v(A_1)}{v(A'_1)} + \frac{v(A_2)}{v(A'_2)} - 2 \\
&= \frac{T' + \delta}{T' + \varepsilon} + \frac{T' - \delta}{T' - \varepsilon} - 2 \\
&= (\delta - \varepsilon) \cdot \left[\frac{1}{T' + \varepsilon} - \frac{1}{T' - \varepsilon} \right] < 0,
\end{aligned}$$

which contradicts the fact that $\vec{A}$ is an allocation satisfying $\Gamma(1)$. Hence, $\vec{A}$ must produce a minimum-difference partition of S'.

To solve the original PARTITION instance, we first compute an allocation satisfying $\Gamma(1)$, use it to produce a minimum-difference partition of S', and then check if its difference is 1. ■

Thus, proportional fairness and the MNW solution, which coincide for allocation of divisible goods, are connected on a spectrum in the case of indivisible goods. Proportional fairness is now strictly stronger, but, unlike the MNW solution, not guaranteed to exist. The spectrum allows us to view the MNW solution as the optimal solution that is guaranteed to exist. Further, among all the solutions on the spectrum that are guaranteed to exist, it is optimally fair, and yet not qualitatively harder in terms of computational complexity.

An interesting potential application of the spectrum framework is to break ties among the set of all MNW allocations. In particular, given an instance of the fair division problem, we can compute the minimum p for which an allocation satisfying $\Gamma(p)$ exists, and compute such an allocation. This approach is guaranteed to select an MNW allocation (Theorem A.3), and can be viewed as the optimal tie-breaking rule. Needless to say, the key challenge will be to develop a scalable implementation of this approach.

Appendix B

Omitted Proofs and Results for Chapter 7

B.1 The MLE rule may not always have the optimal sample complexity

Here, we demonstrate via an example that the MLE rule for a noise model need not always be the rule with the optimal sample complexity in general.

Example B.1. Consider a scenario where there are 3 possible underlying ground truths — σ_1, σ_2 and σ_3. These map to underlying true ranking in the voting context. Let there be 4 possible outcomes — π_1 through π_4. The outcomes map to samples from Mallows' model in our voting context. In the table below, the entry in row i and column j gives the probability of observing outcome π_j given that the ground truth is σ_i, i.e., $\Pr[\pi_j|\sigma_i]$.

	π_1	π_2	π_3	π_4
σ_1	$^1/_5$	$^1/_5$	$^1/_5$	$^2/_5$
σ_2	$^1/_6$	$^1/_6$	$^1/_6$	$^1/_2$
σ_3	$^1/_4$	$^1/_4$	$^1/_4$	$^1/_4$

Take $\varepsilon = 4/5$, so the accuracy requirement is $1 - \varepsilon = 1/5$. Given just one sample from the noise model, the circled entries in the table show the ground truth returned by the MLE rule for various outcomes. It is clear that the MLE rule never returns σ_1, thus it does not achieve the minimum (over all ground truths) accuracy of $1/5$. In contrast, consider the rule which is identical to the MLE rule except that it returns σ_1 when observing π_3. It is clear that given one sample, this rule returns the ground truth with probability at least $1/5$ no matter what the ground truth is. Hence, the sample complexity of the new rule is strictly less than that of the MLE rule for $\varepsilon = 4/5$. This shows that the MLE rule need not always be optimal in terms of its sample complexity.

B.2 Several classical voting rules are PM-c

In this section, we prove that the Kemeny rule, the Slater rule, the ranked pairs method, Copeland's method and Schulze's method are PM-c (Theorem 7.2). The definition of the Kemeny rule is given in Section 7.2. We define the remaining methods below.

The Slater rule Given any profile π, the Slater rule selects a ranking that minimizes the number of pairs of alternatives on which it disagrees with the PM graph of π. Note that this is very similar to the Kemeny rule which, instead of minimizing the number of pairwise disagreements with the PM graph, minimizes the weighted pairwise disagreements, where the weight of a pair is the number of votes by which the stronger alternative beats the weaker alternative.

Copeland's method We say that alternative a beats alternative a' in a profile π if $|\{\sigma \in \pi | a \succ_\sigma a'\}| > |\{\sigma \in \pi | a' \succ_\sigma a\}|$, i.e., if there is an edge from a to a' in the PM graph of π. The Copeland score of an alternative a is the number of alternatives it beats in π and Copeland's method ranks the alternatives in the non-increasing order of their Copeland scores.

The ranked pairs method Under the ranked pairs method, all ordered pairs of alternatives (a, a') are sorted by the number of rankings in the profile in which alternative a is preferred to a' (in non-increasing order). Then, starting with the first pair in the list, the method "locks in" the outcome with the result of the pairwise comparison. It proceeds with the next pairs and locks in every pairwise result that does not contradict the partial ordering established so far (by forming a cycle). Finally, the method outputs the total order obtained.

Schulze's method Given any profile π, define $w(a, a') = |\{\sigma \in \pi | a \succ_\sigma a'\}|$ for any $a, a' \in A$. Consider the (directed) weighted pairwise comparison graph $G = (V, E)$ where the alternatives are the vertices ($V = A$) and there is an edge from every $a \in A$ to every other $a' \in A$ with weight $w(a, a')$. A path of strength t from $a \in A$ to $a' \in A$ is a sequence of vertices $v_0 = a, v_1, \ldots, v_{k-1}, v_k = a'$ where $w(v_i, v_{i+1}) \geqslant t$ for all $i \in \{0, \ldots, k-1\}$, and $w(v_i, v_{i+1}) = t$ for some $i \in \{0, \ldots, k-1\}$. Define the strength of alternative a over alternative a', denoted $s[a, a']$, to be the maximum strength of any path from a to a'. Schulze's method ranks $a \succ a'$ if $s[a, a'] > s[a', a]$. A tie-breaking scheme is used when $s[a, a'] = s[a', a]$.

of Theorem 7.2. Take any ranking $\sigma^* = (a_1 \succ \ldots \succ a_m) \in \mathcal{L}(A)$. We show that for each of the four rules, whenever the PM graph of a profile reduces to σ^*, the rule outputs σ^* with probability 1. Consider any profile π with n votes such that its PM graph reduces to σ^*.

First, note that the Kemeny rule returns a ranking that minimizes the total pairwise disagreements with the input votes. If we consider the (directed) weighted pairwise comparison graph described in the definition of Schulze's method above, then the Kemeny score of a ranking σ, denoted $KemSc(\sigma)$, measures the total pairwise disagreements of σ with the input votes, i.e., $KemSc(\sigma) = \sum_{a,a' \in A | a \succ_\sigma a'} w(a', a)$. In this summation, exactly one edge of the two edges between any pair of alternatives is added. In our profile π, for any $a, a' \in A$ with $a \succ_\sigma a'$, $w(a, a') > n/2 > w(a', a)$. Hence, $KemSc(\sigma^*)$ adds the smaller-weight edge of the two edges between any pair of alternatives. Therefore, σ^* has the minimum Kemeny score, and the Kemeny rule returns σ^*.

The argument for the Slater rule is even easier. Since the PM graph of π reduces to σ^*, σ^* agrees with the PM graph on all pairs of alternatives (and disagrees on none). Therefore, the Slater ranking of π, the ranking that minimizes the number of disagreements with the PM graph on pairs of alternatives, is clearly σ^*. Interestingly, the fact that the Slater rule is PM-c also follows from the distance rationalization of the rule according to the transitivity consensus [143].

When the ranked pairs method is applied to a profile that reduces to σ^*, every ordered pair (a, a') with $a \succ_{\sigma^*} a'$ will be placed before every ordered pair (b, b') with $b' \succ_{\sigma^*} b$. This is because the former pair would be consistent with more than half of the rankings in the profile, while the latter pair would be consistent with less than half of the rankings in the profile. Hence, the ranked pairs method would lock every pair (a, a') where $a \succ_{\sigma^*} a'$ (and obtain the total order σ^*) before reaching any pair of the opposite direction. Therefore, the ranked pairs method would also output σ^*.

For Copeland's method, note that when the PM graph reduces to σ^*, then alternative a_i has Copeland score $m - i$, for every $i \in \{1, \ldots, m\}$. Therefore, Copeland's method outputs exactly the ranking σ^*.

Finally, for Schulze's method, note that for any $a, a' \in A$ with $a \succ_{\sigma^*} a'$, $s[a, a'] > n/2$ (because the edge a to a' itself has weight more than $n/2$). On the other hand, $s[a', a] \leqslant n/2$, because otherwise there would be a cycle in the PM graph that consists of the strongest path from a' to a and the edge from a to a'. Hence, Schulze's method ranks $a \succ a'$ for every $a, a' \in A$ with $a \succ_{\sigma^*} a'$. We conclude that Schulze's method also outputs σ^*.

We have thus established that the Kemeny rule, the Slater rule, the ranked pairs method, Copeland's method and Schulze's method are all PM-c. ■ (Theorem 7.2)

B.3 Several classical voting rules are PD-c

We first define the Bucklin rule and then prove Theorem 7.7.

The Bucklin rule The Bucklin score of an alternative a is the minimum k such that a is among the first k positions in the majority of input votes. The Bucklin rule sorts the alternatives in non-decreasing order according to their Bucklin score, and breaks ties among alternatives with the same Bucklin score ℓ by the number of rankings that have the alternative in the first ℓ positions, with the remaining ties broken arbitrarily.

of Theorem 7.7. Consider a profile π with n rankings such that its PD graph reduces to the ranking σ^*, and let a_i denote the alternative at position i in σ^*. We show that any positional scoring rule as well as the Bucklin rule outputs σ^* on π.

For the Bucklin rule, consider any two alternatives $a, a' \in A$ such that $a \succ_{\sigma^*} a'$. For any $j \in \{1, \ldots, m-1\}$, let $s_j(c)$ denote the number of votes where alternative c is among the first j positions. Let k denote the Bucklin score of a and k' denote the Bucklin score of a'. If $k > k'$, then $s_{k'}(a') > n/2$ and $s_{k'}(a) \leqslant s_{k-1}(a) \leqslant n/2$, which is impossible since the PD graph reduces to σ^*. If $k < k'$, then the Bucklin rule ranks $a \succ a'$, as required. If $k = k'$ and $k \neq m$, then again since the PD graph reduces to σ^*, we have that $s_k(a) > s_k(a')$, so the tie is broken in favor of a. Lastly, we note that $k = k' = m$ is not possible since then it would imply that the total number of appearances of a and a' in the last position is $n - s_{m-1}(a) + n - s_{m-1}(a') > 2 \cdot n - 2 \cdot s_{m-1}(a) \geqslant n$. Thus, for every $a \succ_{\sigma^*} a'$, the Bucklin rule ranks a above a'. We conclude that the Bucklin rule outputs σ^*.

For positional scoring rules, we can follow the reasoning of the proof of Theorem 7.6 and express the score of alternative a_i as $\sum_{j=1}^{m-1}(\beta_j \cdot s_j(a_i)) + n\alpha_m$. Then, the desired fact that the score of a_i is higher than that of a_k when $1 \leqslant i < j \leqslant m$ holds since $s_j(a_i) > s_j(a_k)$ for every $j \in \{1, ..., m-1\}$. ■ (Theorem 7.7)

B.4 Distance Functions

In this section, we prove various properties of the three popular distance functions studied in Chapter 7.

B.4.1 All three of our popular distance functions are both MC and PC

First we give a proof of Lemma 7.7, showing that any swap-increasing distance is both MC and PD. Lemma 7.3 would then imply that the KT distance is both MC and PC.

of Lemma 7.7. In Section 7.5.2, we already argued that any swap-increasing distance is MC. Take any $\sigma^* \in \mathcal{L}(A)$ and $a, b \in A$ such that $a \succ_{\sigma^*} b$. The mapping from every σ to $\sigma_{a \leftrightarrow b}$ is a bijection from $\mathcal{L}_{a \succ b}(A)$ to $\mathcal{L}_{b \succ a}(A)$ and it increases the distance by at least 1 (since any ranking in the domain $\mathcal{L}_{a \succ b}(A)$ follows $a \succ b$). Hence, the mapping is weakly-distance-increasing with respect to σ^*. It follows from Lemma 7.5 that any swap-increasing distance is MC.

To show that it is also PC, fix any $\sigma^* \in \mathcal{L}(A)$, $a, b \in A$ such that $a \succ_{\sigma^*} b$ and $j \in \{1, \ldots, m-1\}$. We wish to show that there exists a bijection $f : \mathcal{S}_j(a) \to \mathcal{S}_j(b)$ which is distance-increasing. Note that we cannot use the mapping from every σ to $\sigma_{a \leftrightarrow b}$ as before, since not every ranking in the domain $\mathcal{S}_j(a)$ follows $a \succ b$ and therefore such a mapping would not be guaranteed to increase distance. Instead, we decompose the domain into three parts: $T = \mathcal{S}_j(a) \cap \mathcal{S}_j(b)$, $D_1 = \{\sigma \in \mathcal{S}_j(a) | \sigma(b) > j\}$, and $D_2 = \{\sigma \in \mathcal{S}_j(b) | \sigma(a) > j\}$. Therefore, $\mathcal{S}_j(a) = T \cup D_1$ and $\mathcal{S}_j(b) = T \cup D_2$.

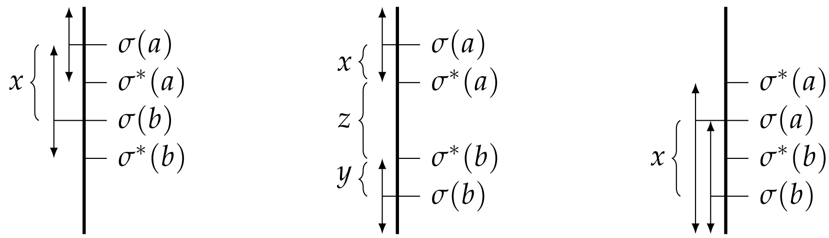

(a) Case 1 (b) Case 2 (c) Case 3

Figure B.1: Exchanges under the footrule and the maximum displacement distances.

Consider the *identity* bijection $I : T \to T$ which maps every ranking to itself. Clearly, I is weakly-distance-increasing with respect to σ^* since it does not change the distance of any ranking from σ^*. Note that for any $\sigma \in D_1$, $\sigma(a) \leqslant j$ and $\sigma(b) > j$, so $a \succ_\sigma b$. Further, $\sigma_{a\leftrightarrow b}(a) > j$ and $\sigma_{a\leftrightarrow b}(b) \leqslant j$. Thus, $\sigma_{a\leftrightarrow b} \in D_2$ and $d(\sigma_{a\leftrightarrow b}, \sigma^*) \geqslant d(\sigma, \sigma^*) + 1$ (by definition). Therefore, the mapping $E : D_1 \to D_2$ where $E(\sigma) = \sigma_{a\leftrightarrow b}$ is distance-increasing with respect to σ^*. Combining the two, the joint bijection $F : \mathcal{S}_j(a) \to \mathcal{S}_j(b)$ naturally given by $F(\sigma) = I(\sigma)$ when $\sigma \in T$ and $F(\sigma) = E(\sigma)$ when $\sigma \in D_1$ is weakly-distance-increasing with respect to σ^*. Further, it is easy to verify that $D_1 \neq \emptyset$, and F increases the distance on any ranking from D_1. Therefore, F is distance-increasing, as required. ■ (Lemma 7.7)

of Theorem 7.12. Lemma 7.7 and Lemma 7.3 already imply that the KT distance is both MC and PC. Now we show that the same holds for the footrule distance (d_{FR}) and the maximum displacement distance (d_{MD}) as well.

Fix any $\sigma^* \in \mathcal{L}(A)$ and $a, b \in A$ such that $a \succ_{\sigma^*} b$. First, we show that for both d_{FR} and d_{MD}, the mapping from every ranking σ with $a \succ_\sigma b$ to $\sigma_{a\leftrightarrow b}$ is weakly-distance-increasing with respect to σ^*. Fix any ranking σ with $a \succ_\sigma b$. We have $\sigma^*(a) < \sigma^*(b)$ and $\sigma(a) < \sigma(b)$. Let $\sigma' = \sigma_{a\leftrightarrow b}$. Recall that $\sigma'(a) = \sigma(b)$ and $\sigma'(b) = \sigma(a)$. For any $c \in A$, let $f(c) = |\sigma(c) - \sigma^*(c)|$ and $f'(c) = |\sigma'(c) - \sigma^*(c)|$ be the displacements of c in σ and σ' respectively. Therefore, $d_{FR}(\sigma, \sigma^*) = \sum_{c\in A} f(c)$, $d_{FR}(\sigma', \sigma^*) = \sum_{c\in A} f'(c)$, $d_{MD}(\sigma, \sigma^*) = \max_{c\in A} f(c)$, $d_{MD}(\sigma', \sigma^*) = \max_{c\in A} f'(c)$. We want to show that $d_{FR}(\sigma', \sigma^*) \geqslant d_{FR}(\sigma, \sigma^*)$ and $d_{MD}(\sigma', \sigma^*) \geqslant d_{MD}(\sigma, \sigma^*)$. Note that $f(c) = f'(c)$ for any $c \in A \setminus \{a, b\}$ since exchanging a and b does not change the positions of the other alternatives. Thus, for the footrule distance it is sufficient to show that $f'(a) + f'(b) \geqslant f(a) + f(b)$, and for the maximum displacement distance it is sufficient to show that $\max(f'(a), f'(b)) \geqslant \max(f(a), f(b))$. We consider three cases.

Case 1. Let $\sigma(a) \leqslant \sigma^*(a)$ and $\sigma(a) < \sigma(b) \leqslant \sigma^*(b)$ as shown in Figure B.1(a). Let $x = \sigma(b) - \sigma(a)$. From the figure, it is easy to verify that by exchanging a and b in σ, b moves farther from $\sigma^*(b)$ by exactly x and a may move closer to $\sigma^*(a)$ but by at most x. Formally,

$$f'(b) - f(b) = (\sigma^*(b) - \sigma'(b)) - (\sigma^*(b) - \sigma(b)) = \sigma(b) - \sigma(a) = x, \tag{B.1}$$

where the second transition holds because $\sigma'(b) = \sigma(a)$. Similarly,

$$f'(a) - f(a) = |\sigma'(a) - \sigma^*(a)| - |\sigma^*(a) - \sigma(a)| \geqslant -|\sigma'(a) - \sigma(a)| = -(\sigma(b) - \sigma(a)) = -x, \tag{B.2}$$

where the second transition is due to triangle inequality and the third transition holds because $\sigma'(a) = \sigma(b)$. Adding Equations (B.1) and (B.2), we get that $f'(a) + f'(b) \geqslant f(a) + f(b)$. For the maximum displacement distance, note that the displacement of b in σ' is $\sigma^*(b) - \sigma(a)$, which is clearly at least as much as the displacements of a and b in σ. Formally,

$$\max(f'(a), f'(b)) = \max(|\sigma(b) - \sigma^*(a)|, \sigma^*(b) - \sigma(a)) = \sigma^*(b) - \sigma(a),$$

where the second transition holds because $\sigma(a) < \sigma(b) \leqslant \sigma^*(b)$ and $\sigma(a) \leqslant \sigma^*(a) < \sigma^*(b)$. Also, $\sigma^*(b) > \sigma^*(a)$ implies $\sigma^*(b) - \sigma(a) > \sigma^*(a) - \sigma(a) = f(a)$ and $\sigma(a) < \sigma(b)$ implies $\sigma^*(b) - \sigma(a) > \sigma^*(b) - \sigma(b) = f(b)$. Hence, $\max(f'(a), f'(b)) > f(a)$ and $\max(f'(a), f'(b)) > f(b)$, so $\max(f'(a), f'(b)) > \max(f(a), f(b))$, as required.

Case 2. Let $\sigma(a) \leqslant \sigma^*(a)$ and $\sigma^*(b) < \sigma(b)$ as shown in Figure B.1(b). Let $x = \sigma^*(a) - \sigma(a)$, $y = \sigma(b) - \sigma^*(b)$ and $z = \sigma^*(b) - \sigma^*(a)$. Then, it is clear that $f(a) = x$, $f(b) = y$, $f'(a) = z + y$, and $f'(b) = z + x$. It is trivial to check that $f'(a) + f'(b) \geqslant f(a) + f(b)$ and $\max(f'(a), f'(b)) \geqslant \max(f(a), f(b))$.

Case 3. Let $\sigma(a) \geqslant \sigma^*(a)$ and $\sigma(b) > \sigma(a)$ as shown in Figure B.1(c). This case is very similar to Case 1. For the footrule distance, alternative a (rather than b) moves away by exactly $x = \sigma(b) - \sigma(a)$ and alternative b (rather than a) may move closer by at most x. Similarly, for the maximum displacement distance, alternative a (rather than b) has greater displacement after the exchange compared to the displacements of both alternatives before the exchange. Hence, we again have $f'(a) + f'(b) \geqslant f(a) + f(b)$ and $\max(f'(a), f'(b)) \geqslant \max(f(a), f(b))$.

From the above three cases, it follows that the mapping which exchanges a and b in a ranking σ with $a \succ_\sigma b$ is weakly-distance-increasing with respect to σ^* for both the footrule distance and the maximum displacement distance. Similarly to the proof of Theorem 7.12, this mapping is a weakly-distance-increasing bijection from $\mathcal{L}_{a \succ b}(A)$ to $\mathcal{L}_{b \succ a}(A)$, which shows that both distances are MC (using the equivalent representation of MC distances given in Lemma 7.5).

To prove that both distances are PC, we use the same technique that we used in the proof of Theorem 7.12. We want to give a bijection from $\mathcal{S}_j(a)$ to $\mathcal{S}_j(b)$ that is distance-increasing. We map every ranking where both a and b are in the first j positions to itself, which does not change the distance of the ranking from σ^*. We map any ranking σ where $\sigma(a) \leqslant j$ and $\sigma(b) > j$ to the ranking where alternatives a and b are swapped, which does not decrease the distance from σ^* as shown in the three cases above. Therefore, this mapping is at least weakly-distance-increasing. We need to show that it is distance-increasing. That is, the distance must increase for some $\sigma \in \mathcal{S}_j(a)$. Clearly, the identity map does not change the distance. Thus, it is sufficient to show that for any $\sigma^* \in \mathcal{L}(A)$, $a, b \in A$ such that $a \succ_{\sigma^*} b$ and $j \in \{1, \ldots, m-1\}$, there exists a ranking σ

such that $\sigma(a) \leqslant j$, $\sigma(b) > j$ and $d(\sigma_{a\leftrightarrow b}, \sigma^*) > d(\sigma, \sigma^*)$ for both $d = d_{FR}$ and $d = d_{MD}$ (note that we can in principle show different rankings for d_{FR} and d_{MD}, but we give a stronger example that works for both distances). For this, we again consider two cases. Let a_1 and a_m denote the first and the last alternatives in σ^*.

Case 1. If $1 \leqslant j < \sigma^*(b)$, then consider the ranking σ where $\sigma(a) = \sigma^*(a_1) = 1$, $\sigma(a_1) = \sigma^*(a)$ and $\sigma(c) = \sigma^*(c)$ for every $c \in A \setminus \{a_1, a\}$. In particular, $\sigma(b) = \sigma^*(b)$. First, note that $\sigma(a) = 1 \leqslant j$ and $\sigma(b) = \sigma^*(b) > j$. Now, it is easy to verify that $f(a) = \sigma^*(a) - 1$, $f(b) = 0$, $f'(a) = \sigma^*(b) - \sigma^*(a)$, and $f'(b) = \sigma^*(b) - 1$. Therefore,

$$\begin{aligned} f'(a) + f'(b) &= 2 \cdot \sigma^*(b) - \sigma^*(a) - 1 \geqslant 2 \cdot (\sigma^*(a) + 1) - \sigma^*(a) - 1 \\ &= \sigma^*(a) + 1 > \sigma^*(a) - 1 = f(a) + f(b). \end{aligned}$$

Therefore, the footrule distance strictly increases. For the maximum displacement distance, note that in the original ranking, only alternatives a and a_1 are displaced, hence $d_{MD}(\sigma, \sigma^*) = f(a)$. Also, in the final ranking, only alternatives a_1, a and b are displaced, and b has the highest displacement. Thus, $d_{MD}(\sigma', \sigma^*) = f'(b)$. Finally, note that $f'(b) = \sigma^*(b) - 1 > \sigma^*(a) - 1 = f(a)$. We conclude that the maximum displacement distance also strictly increases.

Case 2. If $\sigma^*(b) \leqslant j < m$, then consider the ranking σ where $\sigma(b) = \sigma^*(a_m) = m$, $\sigma(a_m) = \sigma^*(b)$ and $\sigma(c) = \sigma^*(c)$ for every $c \in A \setminus \{a_m, b\}$. In particular, $\sigma(a) = \sigma^*(a)$. Again note that $\sigma(a) = \sigma^*(a) < \sigma^*(b) \leqslant j$ and $\sigma(b) = m > j$. Now, it is easy to verify that $f(a) = 0$, $f(b) = m - \sigma^*(b)$, $f'(a) = m - \sigma^*(a)$, and $f'(b) = \sigma^*(b) - \sigma^*(a)$. Therefore,

$$f'(a) + f'(b) = m + \sigma^*(b) - 2 \cdot \sigma^*(a) > m + \sigma^*(b) - 2 \cdot \sigma^*(b) = m - \sigma^*(b) = f(a) + f(b).$$

It follows that the footrule distance strictly increases. For the maximum displacement distance, note that in the original ranking, only alternatives b and a_m are displaced, hence $d_{MD}(\sigma, \sigma^*) = f(b)$. Also, in the final ranking, only alternatives a_m, a and b are displaced, and a has the highest displacement. Thus, $d_{MD}(\sigma', \sigma^*) = f'(a)$. Finally, note that $f'(a) = m - \sigma^*(a) > m - \sigma^*(b) = f(b)$. Hence, the maximum displacement distance also strictly increases.

From both cases, it is clear that for any $\sigma^* \in \mathcal{L}(A)$, $a, b \in A$ such that $a \succ_{\sigma^*} b$ and $j \in \{1, \ldots, m-1\}$, the bijection we constructed from $\mathcal{S}_j(a)$ to $\mathcal{S}_j(b)$ is distance-increasing. Hence, using the equivalent representation of PC distances given in Lemma 7.6, it follows that both the footrule distance and the maximum displacement distance are PC, as required. ■ (Theorem 7.12)

B.4.2 The curious case of the Cayley distance and the Hamming distance

In the discussion, we mentioned that we exclude distances such as the Cayley distance and the Hamming distance from our analysis because even the most prominent voting rules such as plurality are not accurate in the limit for any noise models that are monotonic with respect to these distances. We show that this is indeed the case. First, let us define these two distances.

The Cayley Distance The Cayley distance between two rankings measures the minimum number of (possibly non-adjacent) swaps of alternatives required to convert one ranking into the other. Let us denote it by d_{CY}.

The Hamming Distance The hamming distance between two rankings is defined as the number of positions in which the rankings differ. Formally, $d_{HM}(\sigma_1, \sigma_2) = \sum_{a \in A} \mathbb{1}[\sigma_1(a) \neq \sigma_2(a)]$.

Let $A = \{a, b, c\}$ be the set of alternatives. Let $\sigma^* = (a \succ b \succ c)$ be the true ranking. The following table describes the various possible rankings over these three alternatives and their Cayley distances as well as Hamming distances from σ^*.

	σ_1	σ_2	σ_3	σ_4	σ_5	σ_6
	a	a	b	c	b	c
	b	c	a	b	c	a
	c	b	c	a	a	b
d_{HM}	0		2		3	
d_{CY}	0		1		2	

Recall that for any d-monotonic noise model, $d(\sigma, \sigma^*) = d(\tau, \sigma^*)$ implies $\Pr[\sigma|\sigma^*] = \Pr[\tau|\sigma^*]$. Therefore, any noise model that is monotonic with respect to the Hamming distance or the Cayley distance would assign equal probabilities to rankings σ_3 and σ_4, and to rankings σ_5 and σ_6, making the probabilities of alternatives b and c appearing first in a random vote equal. It follows that with probability $1/2$, alternative c would be ranked higher than alternative b by plurality. Thus, plurality is not accurate in the limit with respect to any noise model that is monotonic with respect to either the Hamming distance or the Cayley distance.

B.5 Two Useful Lemmas

Lemma B.1 (Convexity Lemma). *Consider a point $x \in \Delta^{m!}$. Let $FIX \subseteq \mathcal{L}(A)$, and $VARY = \mathcal{L}(A) \setminus FIX$. Further, assume that $\{\sigma \in \mathcal{L}(A) | x_\sigma = 0\} \subseteq FIX$, and let $k = |VARY| \geqslant 2$. Define*

$$V = \Big\{ v \in \{-1, 0, 1\}^{m!} \;\Big|\; \forall \sigma \in \mathcal{L}(A),\, v_\sigma = 0 \Leftrightarrow \sigma \in FIX \\ \wedge\, \exists \sigma \in \mathcal{L}(A),\, v_\sigma = 1 \wedge \exists \sigma \in \mathcal{L}(A),\, v_\sigma = -1 \Big\}.$$

For every $v \in V$, define the orthant

$$O^v = \Big\{ y \in \Delta^{m!} \;\Big|\; \forall \sigma \in \mathcal{L}(A), (v_\sigma = 0 \Rightarrow y_\sigma = x_\sigma) \wedge (v_\sigma = 1 \Rightarrow y_\sigma > x_\sigma)$$

$$\wedge (v_\sigma = -1 \Rightarrow y_\sigma < x_\sigma) \Big\}.$$

Given points $x^v \in O^v$ for all $v \in V$, $x \in \mathbf{co}\{x^v | v \in V\}$, where $\mathbf{co}$ denotes the convex hull.

Proof. We prove this by induction on k.

For $k = 2$, let $VARY = \{\sigma_1, \sigma_2\}$. Thus, O^v contains two orthants: one consisting of y's where $y_{\sigma_1} < x_{\sigma_1}$ and $y_{\sigma_2} > x_{\sigma_2}$, and another consisting of y's where $y_{\sigma_1} > x_{\sigma_1}$ and $y_{\sigma_2} < x_{\sigma_2}$. We are given a point x^1 in the former orthant and a point x^2 in the latter orthant. For both points, the values of coordinates other than σ_1 and σ_2 match those for x. Hence, it is easy to check that $x = \lambda x^1 + (1-\lambda)x^2$, where $\lambda = (x_{\sigma_1} - x^2_{\sigma_1})/(x^1_{\sigma_1} - x^2_{\sigma_1})$. It is further easy to check that $0 < \lambda < 1$. Hence, $x \in \mathsf{co}\{x^1, x^2\}$.

Suppose that the theorem holds for all $FIX, VARY$ with $k = |VARY| = d - 1$, for some $d \leqslant m!$. Let us consider $FIX, VARY$ with $k = |VARY| = d$. Define V and O^v for every $v \in V$ from FIX. Take any $\tau \in VARY$, construct $\widehat{FIX} = FIX \cup \{\tau\}$, and define $\widehat{V}$ and $\widehat{O^v}$ for every $v \in \widehat{V}$ according to $\widehat{FIX}$.

If we can find a point $\widehat{x^v} \in \widehat{O^v}$ for each $v \in \widehat{V}$ which is also in $\mathsf{co}\{x^v | v \in V\}$, then by the induction hypothesis, we have $x \in \mathsf{co}\{\widehat{x^v} | v \in \widehat{V}\} \subseteq \mathsf{co}\{x^v | v \in V\}$. Take any $v \in \widehat{V}$. We construct $v^{+1}, v^{-1} \in V$ as follows: $v^+_\tau = +$, $v^-_\tau = -$, and $v^+_\sigma v^-_\sigma = v_\sigma$ for all $\sigma \neq \tau$. We show that we can find $\widehat{x^v} \in \widehat{O^v}$ as a convex combination of x^{v^+} and x^{v^-}. It is easy to check that taking $\widehat{x^v} = \lambda x^{v^+} + (1-\lambda)x^{v^-}$ works, where $\lambda = (x_\tau - x^{v^-}_\tau)/(x^{v^+}_\tau - x^{v^-}_\tau)$. It is easy to check that $0 < \lambda < 1$. Further, we have $\widehat{x^v}_\tau = x_\tau$ by construction, which is desired because $\tau \in \widehat{FIX}$. For every $\sigma \neq \tau$, $v^+_\sigma v^-_\sigma = v_\sigma$. Hence,

$$v_\sigma = + \Rightarrow \left(x^{v^+}_\sigma > x_\sigma \wedge x^{v^-}_\sigma > x_\sigma \right) \Rightarrow \widehat{x^v}_\sigma > x_\sigma,$$

and

$$v_\sigma = - \Rightarrow \left(x^{v^+}_\sigma < x_\sigma \wedge x^{v^-}_\sigma < x_\sigma \right) \Rightarrow \widehat{x^v}_\sigma < x_\sigma.$$

For arbitrary $v \in \widehat{V}$, we found $\widehat{x^v} \in \widehat{O^v}$, which is also in $\mathsf{co}\{x^v | v \in V\}$ as desired. Thus, $x \in \mathsf{co}\{x^v | v \in V\}$. ■(Proof of Lemma B.1)

Lemma B.2. *Given a specific ranking $\sigma^* \in \mathcal{L}(A)$ and a probability distribution D over the rankings in $\mathcal{L}(A)$ such that*

$$\arg\max_{\tau \in \mathcal{L}(A)} \Pr_D[\tau] = \{\sigma^*\},$$

there exists a distance metric d over $\mathcal{L}(A)$ and a d-monotonic noise model G with $\Pr_G[\sigma; \sigma^*] = \Pr_D[\sigma]$ for every $\sigma \in \mathcal{L}(A)$.

Proof. First, let $V = \{\Pr_D[\sigma] | \sigma \in \mathcal{L}(A)\}$ be the set of distinct probability values in D. Now, we construct the distance metric d as follows. For all $\sigma \in \mathcal{L}(A)$, set $d(\sigma, \sigma^*) = d(\sigma^*, \sigma) = |\{v \in V | v > \Pr_D[\sigma]\}|$ for every $\sigma \in \mathcal{L}(A)$. For every pair of rankings σ, σ' different than σ^*, we set $d(\sigma, \sigma') = 0$ if $\sigma = \sigma'$ and $d(\sigma, \sigma') = d(\sigma, \sigma^*) + d(\sigma', \sigma^*)$.

We can easily show that the function d is indeed a distance metric. The first two properties are preserved by definition. For the triangle inequality, we wish to prove that $d(\sigma,\sigma') + d(\sigma',\sigma'') \geqslant d(\sigma,\sigma'')$ for all $\sigma,\sigma',\sigma'' \in \mathcal{L}(A)$. The inequality clearly holds when any two of the three rankings are identical. If all three rankings are distinct, we take two cases.

1. Suppose either $\sigma = \sigma^*$ or $\sigma'' = \sigma^*$. Without loss of generality, let us assume $\sigma = \sigma^*$. Then, the above inequality is obvious since, by the definition of d, $d(\sigma',\sigma'') \geqslant d(\sigma^*,\sigma'') = d(\sigma,\sigma'')$.
2. Suppose that neither σ nor σ'' is equal to σ^*. Then, again by the definition of d, the LHS of the above inequality becomes

$$d(\sigma^*,\sigma) + 2d(\sigma^*,\sigma') + d(\sigma^*,\sigma'') \geqslant d(\sigma^*,\sigma) + d(\sigma^*,\sigma'') = d(\sigma,\sigma'').$$

We now define the noise model G as $\Pr_G[\sigma;\sigma^*] = \Pr_D[\sigma]$ for every $\sigma \in \mathcal{L}(A)$ and

$$\Pr_G[\sigma;\sigma'] = \frac{1/(1+d(\sigma,\sigma'))}{\sum_{\tau \in \mathcal{L}(A)} 1/(1+d(\tau,\sigma'))}$$

for $\sigma' \neq \sigma^*$. The property $\Pr_G[\sigma;\sigma'] \geqslant \Pr[\sigma'';\sigma']$ iff $d(\sigma;\sigma') \leqslant d(\sigma'',\sigma')$ is obvious if $\sigma' \neq \sigma^*$. Otherwise, recall that $\Pr_G[\sigma;\sigma^*] = \Pr_D[\sigma]$ and, clearly, $\Pr_G[\sigma;\sigma^*] \geqslant \Pr_G[\sigma';\sigma^*]$ iff $|\{v \in V | v > \Pr_D[\sigma]\}| \leqslant |\{v \in V | v > \Pr_D[\sigma']\}|$, i.e., $d(\sigma,\sigma^*) \leqslant d(\sigma',\sigma^*)$. ■ (Proof of Lemma B.2)

Appendix C

Omitted Proofs and Results for Chapter 8

C.1 Proof of Lemmas

Lemma C.1. *Let σ^* denote the true ranking of the alternatives, and let σ denote a sample ranking obtained from the Thurstone-Mosteller model where all Gaussians have identical variance ν^2, which is upper bounded by a constant M^2. Let $a, b \in A$ be two alternatives such that $a \succ_{\sigma^*} b$. Take a ranking $\tau \in \mathcal{L}(A)$ with $a \succ_\tau b$, and let $\tau_{a \leftrightarrow b}$ denote the ranking obtained by swapping a and b in τ. Then, we have that $\Pr[\sigma = \tau] - \Pr[\sigma = \tau_{a\leftrightarrow b}]$ is non-negative, and is at least a positive constant when $\tau = \sigma^*$.*

Proof. Jiang et al. [115] show (in the proof of their Theorem 2, and under the assumption of a constant upper *and lower* bound on the variance) that for $i < j$, the difference

$$\Omega_{a,b,i,j} = \Pr[\sigma(a) = i \wedge \sigma(b) = j] - \Pr[\sigma(a) = j \wedge \sigma(b) = i]$$

is at least a positive constant. Note that $\Omega_{a,b,i,j}$ can also be expressed as follows. Let $\Omega_\tau = \Pr[\sigma = \tau] - \Pr[\sigma = \tau_{a\leftrightarrow b}]$. Then,

$$\Omega_{a,b,i,j} = \sum_{\tau \in \mathcal{L}(A):\tau(a)=i \wedge \tau(b)=j} \Omega_\tau.$$

Take $i = \sigma^*(a)$ and $j = \sigma^*(b)$. We are interested in the value of Ω_τ for a particular ranking τ, and in the value of $\Omega_{a,b,i,j}$.

Let $\{Y_c\}_{c\in A}$ be the random variables denoting the sampled utilities of the alternatives. Crucially, we observe that Jiang et al. express $\Omega_{a,b,i,j}$ as an integral of a certain density function — which is lower bounded by a positive constant — over the full range of values of $\{Y_c\}_{c\in A\setminus\{a,b\}}$. They show that this integral is over a region with at least a constant probability, thus yielding a positive constant lower bound on $\Omega_{a,b,i,j}$.

Instead of integrating over the full range of values, we integrate only over values of $\{Y_c\}_{c\in A\setminus\{a,b\}}$ that are consistent with the ordering of the alternatives of $A \setminus \{a,b\}$ in τ (or equivalently, in $\tau_{a\leftrightarrow b}$). This integral yields Ω_τ. It is easy to check that the density function (which may not be lower bounded by a positive constant in the absence of a lower bound on the variance) is still positive. Thus, $\Omega_\tau > 0$ for every $\tau \in \mathcal{L}(A)$. However, it is easy to see that for every $\tau \neq \sigma^*$, both $\Pr[\sigma = \tau]$ and $\Pr[\sigma = \tau_{a\leftrightarrow b}]$ go to 0 as ν^2 goes to 0. Hence, Ω_τ also approaches 0 as ν^2 goes to 0.

Finally, in the case of $\tau = \sigma^*$, we want to show that $\inf_{\nu^2 \in (0,M^2]} \Omega_{\sigma^*}$ is lower bounded by a positive constant. Note that when $\nu^2 = M^2$, we know that Ω_{σ^*} is lower bounded by a positive constant due to the result of Jiang et al. [115]. When $\nu^2 = 0$, we have $\Pr[\sigma = \sigma^*] = 1$ and $\Pr[\sigma = \sigma^*_{a\leftrightarrow b}] = 0$. Hence, $\Omega_{\sigma^*} = 1$. By the extreme value theorem, Ω_{σ^*} must achieve its minimum value when $\nu^2 \in [0, M^2]$. However, we already established that $\Omega_\tau > 0$ for every $\tau \in \mathcal{L}(A)$ and every value of ν^2. Hence, this minimum value must be at least a positive constant, as required. ■

Lemma C.2. *Let σ^* denote the true ranking of the alternatives, and let σ denote a sample ranking obtained from the Thurstone-Mosteller model where all Gaussians have identical variance ν^2, which is upper bounded by a constant M^2. Let $a, b \in A$ be two alternatives such that $a \succ_{\sigma^*} b$. Then, we have that $\Pr[\sigma(a) \in [j]] - \Pr[\sigma(b) \in [j]]$ is non-negative for every $j \in \{1, \ldots, m-1\}$ and at least a positive constant for some $j = \sigma^*(a)$.*

Proof. In this case,

$$\Omega_{a,b,j} = \Pr[\sigma(a) \in [j]] - \Pr[\sigma(b) \in [j]],$$

for $a \succ_{\sigma^*} b$ and $j \in \{1, \ldots, m-1\}$. Jiang et al. [115] show (in the proof of their Theorem 2) that $\Omega_{a,b,j}$ is at least a positive constant for every $j \in \{1, \ldots, m-1\}$. However, they use an upper *and a lower* bound on the variance. Clearly, $\Omega_{a,b,j}$ can approach 0 as the variance ν^2 approaches 0, e.g., if $j < \min(\sigma^*(a), \sigma^*(b))$.

As in the proof of Lemma C.1, we observe that Jiang et al. express $\Omega_{a,b,j}$ as an integral of a certain density function — which, in their case, is lower bounded by a positive constant — over a region with probability lower bounded by a positive constant. It is easy to check that even in the absence of a lower bound on the variance, the density function is still positive, yielding $\Omega_{a,b,j} > 0$ for all $a \succ_{\sigma^*} b$ and $j \in \{1, \ldots, m-1\}$.

Further, take $j = \sigma^*(a)$. Due to the results of Jiang et al., we have that $\Omega_{a,b,j}$ is at least a positive constant when $\nu^2 = M^2$. When $\nu^2 = 0$, we have that $\Pr[\sigma(a) \in [j]] = 1$ and $\Pr[\sigma(b) \in [j]] = 0$. Hence, $\Omega_{a,b,j} = 1$. By the extreme value theorem, as ν^2 varies in the interval $[0, M^2]$, $\Omega_{a,b,j}$ achieves its minimum value. Further, since $\Omega_{a,b,j}$ is always positive, this minimum value must be a positive constant, as required. ■

C.2 Plurality Fails With Equal Variance

Let us consider the social network structure which consists of n vertices $\{v_1, \ldots, v_n\}$, where v_1 is only connected to v_2, and $\{v_2, \ldots, v_n\}$ form an $(n-1)$-clique. Vertices v_2 through v_n place equal weights on all their incident edges, and vertex v_1 places weight 1 on its unique incident edge. Note that this satisfies the assumptions of Theorem 8.1.

Next, let the set of alternatives be $A = \{a, b, c\}$, and let their true qualities be $\mu_a = 3$, $\mu_b = 2$, and $\mu_c = 1$. Fix $\nu^2 = 1$. Note that under the distribution where the noisy quality estimate of each alternative is sampled from a Gaussian with its true quality being the mean and variance ν^2, every ranking has a positive probability of being sampled. Let these Gaussians be associated with each edge in the network.

Then, we see that the Gaussians at vertices v_2 through v_n have variance at most $2\nu^2/(n-1)$, which goes to 0 as $n \to \infty$. Hence, in the limit, each of v_2 through v_n would report the true ranking with probability 1. However, vertex v_1 can report each ranking with a constant positive probability. Hence, there is a constant positive probability that vertices v_2 through v_n report the true ranking $(a \succ b \succ c)$, and vertex v_1 reports the ranking $(c \succ a \succ b)$. Thus, plurality would return the ranking $(a \succ c \succ b)$ with a positive probability in the limit as $n \to \infty$, which is violation of accuracy in the limit.

C.3 Borda Count And The Modal Ranking Rule Fail With Unequal Variance

Let there be three alternatives: a, b, and c. Hence, in this example $m = 3$. Let their true qualities be $\mu_a = 9$, $\mu_b = 7$, and $\mu_c = 3$. Hence, the ground truth ranking is $\sigma^* = (a \succ b \succ c)$. Associate a standard deviation with each alternative as follows: $\nu_a = 17$, $\nu_b = 2$, and $\nu_c = 1$. Imagine a noise model where a noisy quality estimate is sampled for each alternative $x \in \{a, b, c\}$ from the Gaussian $\mathcal{N}(\mu_x, (\nu_x)^2)$, and the alternatives are then ranked according to their noisy quality estimates. Let BD_x denote the expected Borda score of alternative $x \in \{a, b, c\}$ in a ranking sampled from this noise model. Then, it can be verified that $BD_b > BD_a > BD_c$.

Next, construct a "star" social network structure among n voters. That is, voter v_1 is connected to every other voter, and there are no other edges in the network. Let the aforementioned noise model be associated with each of the $n-1$ edges in the network. Then, the votes of voters v_2 through v_n are simply $n-1$ independent samples from the aforementioned noise model. Hence, the expected Borda score of b is greater than the expected Borda score of a in each of these $n-1$ votes. It follows that as n goes to infinity, the overall Borda score of b would be greater than the overall Borda score of a with probability 1. Thus, Borda count would not be able to return the ground truth ranking $a \succ b \succ c$ with high probability, violating accuracy in the limit.

Similarly, for the modal ranking rule, let the true qualities be $\mu_a = 9$, $\mu_b = 6$, and $\mu_c = 5$, and the standard deviations be $\nu_a = 18$, $\nu_b = 20$, and $\nu_c = 10$. Then, it can be verified that for a ranking σ sampled from the associated noise model, $\Pr[\sigma = (a \succ c \succ b)] > \Pr[\sigma = (a \succ b \succ c)]$. Once again, for the "star" social network structure described above, voters v_2 through v_n are more likely to submit ranking $(a \succ c \succ b)$ than the true ranking $(a \succ b \succ c)$. Hence, as n goes to infinity, ranking $(a \succ c \succ b)$ would appear more number of times than the true ranking $(a \succ b \succ c)$ with probability 1. Thus, the modal ranking rule would not be able to return the true ranking with high probability, violating accuracy in the limit.

Appendix D

Omitted Proofs and Results for Chapter 9

D.1 Additional Experiments

In Chapter 9, we presented experiments (Figure 9.2) that compare our proposed worst-case optimal rule against other voting rules when: i) it receives the true error of a profile $t^* = d(\pi, \sigma^*)$ as an argument (Figures 9.2(a) and 9.2(b)), and ii) when it receives an estimate $\widehat{t}$ of t^* (Figures 9.2(c) and 9.2(d)). In these experiments, we used the Kendall tau distance as the measure of error. In this section we present additional experiments in an essentially identical setting but using the other three distance metrics considered in this paper as the measure of error. These experiments affirm that our proposed rules are superior to other voting rules independent of the error measure chosen. Figures D.1, D.2, and D.3 show the experiments for the footrule distance, the Cayley distance, and the maximum displacement distance, respectively.

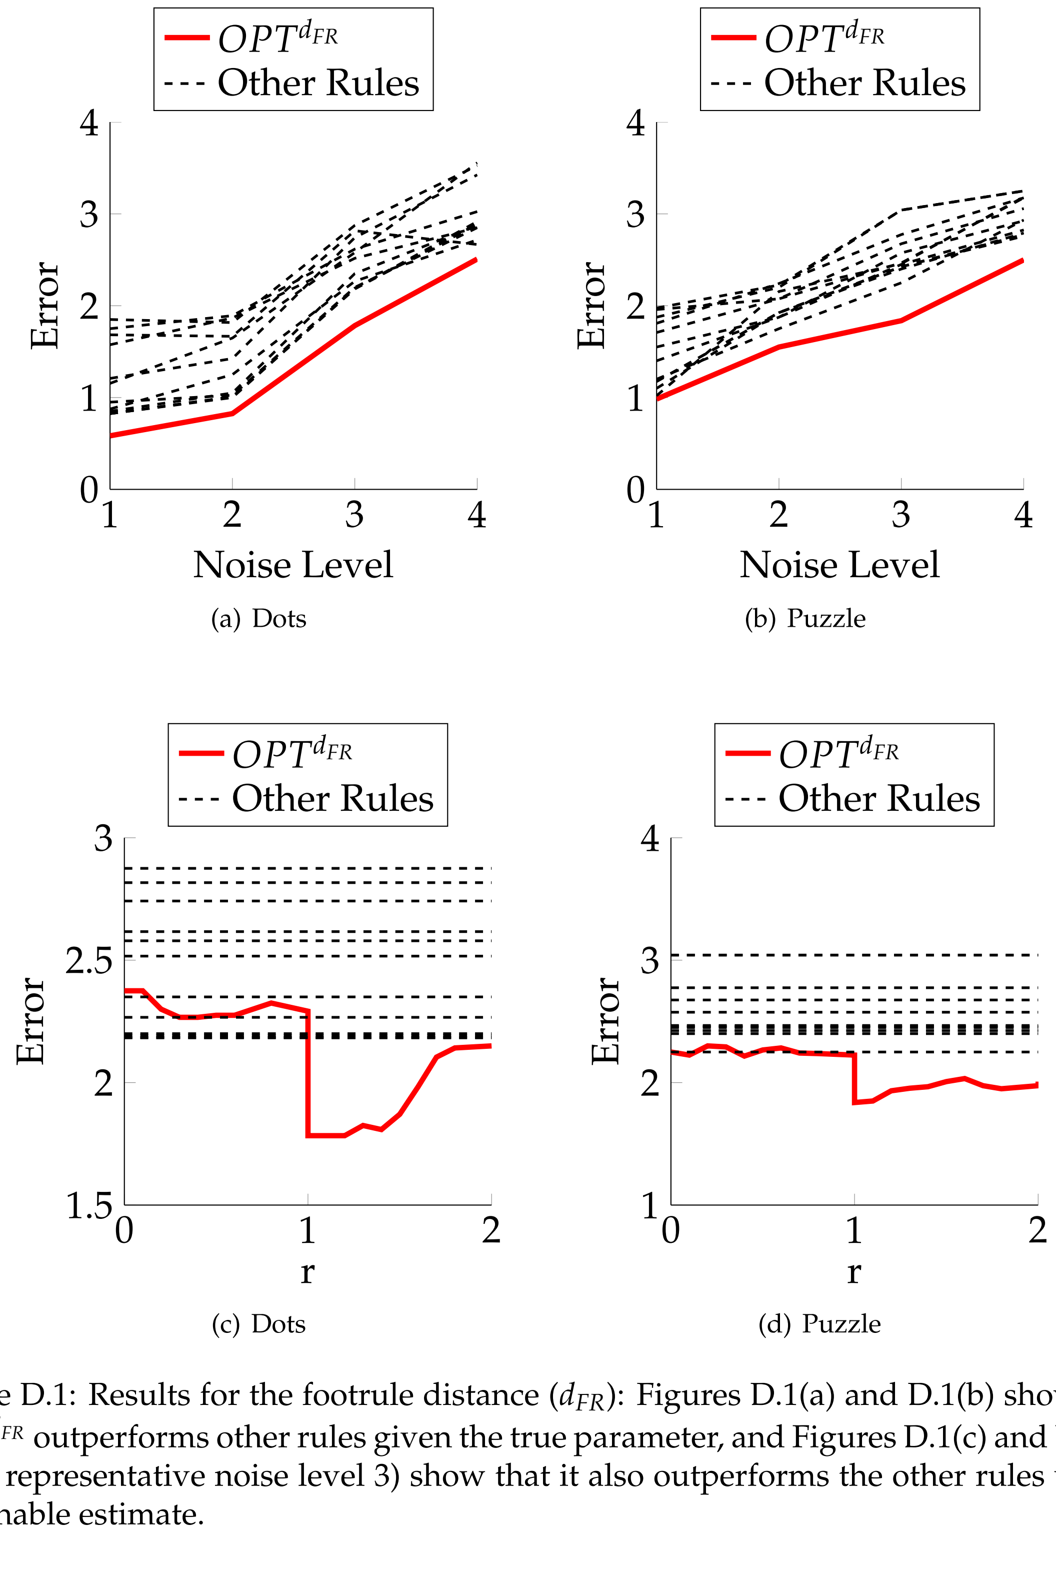

Figure D.1: Results for the footrule distance (d_{FR}): Figures D.1(a) and D.1(b) show that OPT$^{d_{FR}}$ outperforms other rules given the true parameter, and Figures D.1(c) and D.1(d) (for a representative noise level 3) show that it also outperforms the other rules with a reasonable estimate.

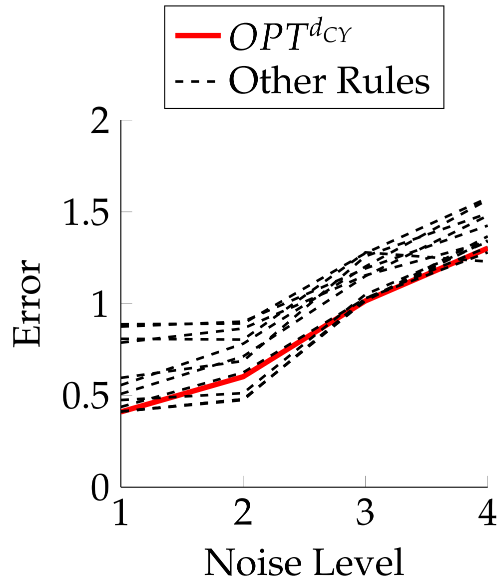

(a) Dots

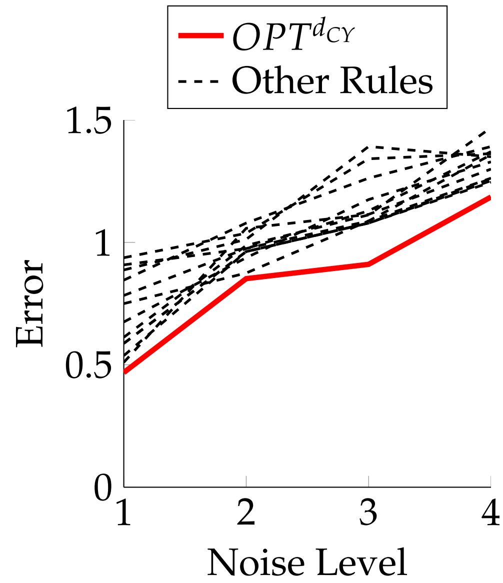

(b) Puzzle

(c) Dots

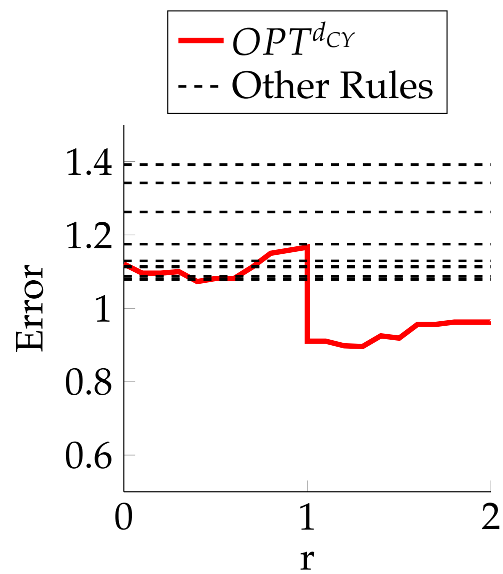

(d) Puzzle

Figure D.2: Results for the Cayley distance (d_{CY}): Figures D.2(a) and D.2(b) show that OPT$^{d_{CY}}$ outperforms other rules given the true parameter, and Figures D.2(c) and D.2(d) (for a representative noise level 3) show that it also outperforms the other rules with a reasonable estimate.

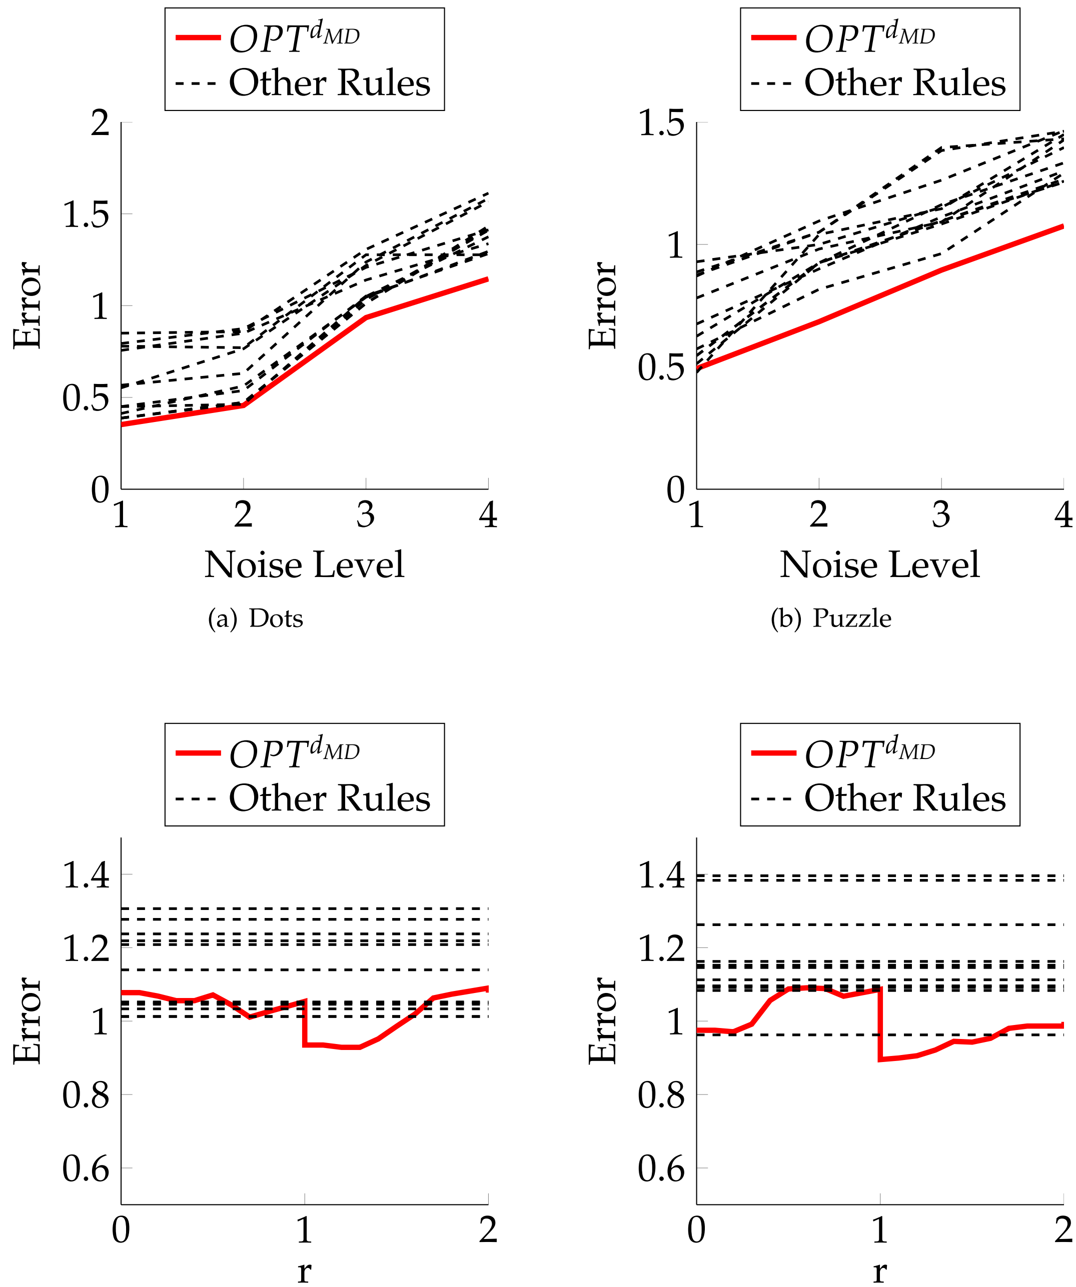

(a) Dots (b) Puzzle

(c) Dots (d) Puzzle

Figure D.3: Results for the maximum displacement distance (d_{MD}): Figures D.3(a) and D.3(b) show that $\text{OPT}^{d_{MD}}$ outperforms other rules given the true parameter, and Figures D.3(c) and D.3(d) (for a representative noise level 3) show that it also outperforms the other rules with a reasonable estimate.

Bibliography

[1] M. Aleksandrov, H. Aziz, S. Gaspers, and T. Walsh. Online fair division: Analysing a food bank problem. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2540–2546, 2015. (page 26)

[2] G. Amanatidis, E. Markakis, A. Nikzad, and A. Saberi. Approximation algorithms for computing maximin share allocations. In *Proceedings of the 42nd International Colloquium on Automata, Languages and Programming (ICALP)*, pages 39–51, 2015. (page 31)

[3] E. Anshelevich and S. Sekar. Blind, greedy, and random: Algorithms for matching and clustering using only ordinal information. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, pages 390–396, 2016. (pages 17 and 142)

[4] E. Anshelevich, O. Bhardwaj, and J. Postl. Approximating optimal social choice under metric preferences. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, pages 777–783, 2015. (pages 17 and 142)

[5] K. Arrow. *Social Choice and Individual Values*. Wiley, 1951. (pages 10 and 117)

[6] K. J. Arrow and M.D. Intriligator, editors. *Handbook of Mathematical Economics*. North-Holland, 1982. (page 25)

[7] H. Azari Soufiani, D. C. Parkes, and L. Xia. Random utility theory for social choice. In *Proceedings of the 26th Annual Conference on Neural Information Processing Systems (NIPS)*, pages 126–134, 2012. (pages 18, 183, 189, and 196)

[8] H. Azari Soufiani, D. C. Parkes, and L. Xia. Preference elicitation for general random utility models. In *Proceedings of the 29th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 596–605, 2013. (page 196)

[9] H. Azari Soufiani, D. C. Parkes, and L. Xia. Computing parametric ranking models via rank-breaking. In *Proceedings of the 31st International Conference on Machine Learning (ICML)*, pages 360–368, 2014. (pages 18 and 196)

[10] H. Aziz. Random assignment with multi-unit demands. arXiv:1401.7700, 2014. (pages 13 and 65)

[11] H. Aziz and S. Mackenzie. A discrete and bounded envy-free cake cutting protocol for any number of agents. arXiv:1604.03655, 2016. (page 8)

[12] H. Aziz, M. Brill, V. Conitzer, E. Elkind, R. Freeman, and T. Walsh. Justified representation in approval-based committee voting. In *Proceedings of the 29th AAAI*

Conference on Artificial Intelligence (AAAI), pages 784–790, 2015. (pages 17 and 142)

[13] H. Aziz, S. Gaspers, S. Mackenzie, and T. Walsh. Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence*, 227:71–92, 2015. (page 41)

[14] Y. Bachrach and N. Shah. Reliability weighted voting games. In *Proceedings of the 6th International Symposium on Algorithmic Game Theory (SAGT)*, pages 38–49, 2013. (page 21)

[15] Y. Bachrach, I. Kash, and N. Shah. Agent failures in totally balanced games and convex games. In *Proceedings of the 8th Conference on Web and Internet Economics (WINE)*, pages 15–29, 2012. (page 21)

[16] Y. Bachrach, R. Savani, and N. Shah. Cooperative max games and agent failures. In *Proceedings of the 13th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 29–36, 2014. (page 21)

[17] A. Barg and A. Mazumdar. Codes in permutations and error correction for rank modulation. *IEEE Transactions on Information Theory*, 56(7):3158–3165, 2010. (pages 20 and 217)

[18] J. Bartholdi and J. Orlin. Single Transferable Vote resists strategic voting. *Social Choice and Welfare*, 8:341–354, 1991. (page 10)

[19] J. Bartholdi, C. A. Tovey, and M. A. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6:227–241, 1989. (pages 10 and 42)

[20] J. Bartholdi, C. A. Tovey, and M. A. Trick. Voting schemes for which it can be difficult to tell who won the election. *Social Choice and Welfare*, 6:157–165, 1989. (pages 10, 183, and 219)

[21] J. Bartholdi, C. A. Tovey, and M. A. Trick. How hard is it to control an election. *Mathematical and Computer Modelling*, 16:27–40, 1992. (page 10)

[22] N. Betzler, M. R. Fellows, J. Guo, R. Niedermeier, and F. A. Rosamond. Fixed-parameter algorithms for Kemeny rankings. *Theoretical Computer Science*, 410(45): 4554–4570, 2009. (page 10)

[23] N. Betzler, J. Guo, C. Komusiewicz, and R. Niedermeier. Average parameterization and partial kernelization for computing medians. *Journal of Computer and System Sciences*, 77(4):774–789, 2011. (page 219)

[24] N. Betzler, R. Bredereck, and R. Niedermeier. Theoretical and empirical evaluation of data reduction for exact Kemeny rank aggregation. *Autonomous Agents and Multi-Agent Systems*, 28(5):721–748, 2014. (page 219)

[25] G. Birkhoff. Three observations on linear algebra. *Universidad Nacional de Tucumán, Revista A*, 5:147–151, 1946. (pages 59, 63, 65, 126, and 131)

[26] A. Blum. *On-line algorithms in machine learning*. Springer, 1998. (pages 20 and 199)

[27] A. Blum and Y. Mansour. Learning, regret minimization, and equilibria. In N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 4. Cambridge University Press, 2007. (pages 4 and 119)

[28] O. Bochet, R. Ilkılıç, H. Moulin, and J. Sethuraman. Balancing supply and demand under bilateral constraints. *Theoretical Economics*, 7(3):395–423, 2012. (pages 4, 9, 13, and 52)

[29] A. Bogomolnaia and H. Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100:295–328, 2001. (pages 13, 41, and 65)

[30] A. Bogomolnaia and H. Moulin. Random matching under dichotomous preferences. *Econometrica*, 72:257–279, 2004. (pages 3, 4, 8, 9, 13, 41, 45, 46, 47, 49, 50, 52, 59, and 63)

[31] A. Bogomolnaia, H. Moulin, and R. Stong. Collective choice under dichotomous preferences. *Journal of Economic Theory*, 122(2):165–184, 2005. (pages 9, 13, 53, and 55)

[32] A. Borodin and R. El-Yaniv. *Online computation and competitive analysis*. Cambridge University Press, 2005. (pages 20 and 199)

[33] C. Boutilier and A. D. Procaccia. A dynamic rationalization of distance rationalizability. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI)*, pages 1278–1284, 2012. (page 183)

[34] C. Boutilier, I. Caragiannis, S. Haber, T. Lu, A. D. Procaccia, and O. Sheffet. Optimal social choice functions: A utilitarian view. In *Proceedings of the 13th ACM Conference on Economics and Computation (EC)*, pages 197–214, 2012. (pages 145 and 218)

[35] C. Boutilier, I. Caragiannis, S. Haber, T. Lu, A. D. Procaccia, and O. Sheffet. Optimal social choice functions: A utilitarian view. *Artificial Intelligence*, 227:190–213, 2015. (pages 4, 11, 17, 117, 118, 119, 125, and 142)

[36] S. Bouveret and M. Lemaître. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Autonomous Agents and Multi-Agent Systems*, 30(2): 259–290, 2016. (page 26)

[37] S. J. Brams and A. D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996. (pages 27 and 41)

[38] S. J. Brams, M. Kilgour, and C. Klamler. Maximin envy-free division of indivisible items. Manuscript, 2015. (page 41)

[39] F. Brandt, V. Conitzer, and U. Endriss. Computational social choice. In G. Weiß, editor, *Multiagent Systems*, chapter 6. MIT Press, 2nd edition, 2013. (page 10)

[40] F. Brandt, V. Conitzer, U. Endress, J. Lang, and A. D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016. (page 17)

[41] M. Braverman and E. Mossel. Noisy sorting without resampling. In *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 268–276, 2008. (page 183)

[42] M. Brill, V. Conitzer, R. Freeman, and N. Shah. False-name-proof recommendations in social networks. In *Proceedings of the 15th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 332–340, 2016.

(page 21)

[43] S. Bubeck and N. Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5(1): 1–122, 2012. (pages 4 and 119)

[44] E. Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011. (pages 3, 7, 8, 12, 13, 26, 40, 65, 82, 86, and 131)

[45] E. Budish, Y.-K. Che, F. Kojima, and P. Milgrom. Designing random allocation mechanisms: Theory and applications. *American Economic Review*, 103(2):585–623, 2013. (pages 13, 41, 65, and 126)

[46] D. E. Campbell and J. S. Kelly. Arrovian social choice correspondences. *International Economic Review*, 37(4):803–823, 1996. (page 119)

[47] I. Caragiannis and A. D. Procaccia. Voting almost maximizes social welfare despite limited communication. *Artificial Intelligence*, 175(9–10):1655–1671, 2011. (pages 17, 118, and 142)

[48] I. Caragiannis, A. D. Procaccia, and N. Shah. When do noisy votes reveal the truth? In *Proceedings of the 14th ACM Conference on Economics and Computation (EC)*, pages 143–160, 2013. (page 11)

[49] I. Caragiannis, A. D. Procaccia, and N. Shah. Modal ranking: A uniquely robust voting rule. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, pages 616–622, 2014. (pages 11 and 21)

[50] I. Caragiannis, D. Kurokawa, H. Moulin, A. D. Procaccia, N. Shah, and J. Wang. The unreasonable fairness of maximum Nash welfare. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, pages 305–322, 2016. (pages 7, 8, 9, and 21)

[51] I. Caragiannis, S. Nath, A. D. Procaccia, and N. Shah. Subset selection via implicit utilitarian voting. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 151–157, 2016. (pages 11 and 21)

[52] I. Caragiannis, A. D. Procaccia, and N. Shah. When do noisy votes reveal the truth? *ACM Transactions on Economics and Computation*, 4(3): article 15, 2016. (page 21)

[53] I. Caragiannis, A. D. Procaccia, and N. Shah. Truthful univariate estimators. In *Proceedings of the 33rd International Conference on Machine Learning (ICML)*, pages 127–135, 2016. (page 21)

[54] J. R. Chamberlin and P. N. Courant. Representative deliberations and representative decisions: Proportional representation and the Borda rule. *American Political Science Review*, 77(3):718–733, 1983. (pages 17, 119, and 142)

[55] Y. K. Che and F. Kojima. Asymptotic equivalence of probabilistic serial and random priority mechanisms. *Econometrica*, 78(5):1625–1672, 2010. (pages 13 and 65)

[56] Y. Chen, J. K. Lai, D. C. Parkes, and A. D. Procaccia. Truth, justice, and cake cutting. *Games and Economic Behavior*, 77:284–297, 2013. (pages 4, 9, 13, and 52)

[57] F. Chierichetti and J. Kleinberg. Voting with limited information and many alternatives. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1036–1055, 2012. (pages 182 and 183)

[58] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971. (page 6)

[59] R. Cole and V. Gkatzelis. Approximating the Nash social welfare with indivisible items. In *Proceedings of the 47th Annual ACM Symposium on Theory of Computing (STOC)*, pages 371–380, 2015. (pages 9, 12, 26, and 40)

[60] Marquis de Condorcet. Essai sur l'application de l'analyse à la probabilité de décisions rendues à la pluralité de voix. Imprimerie Royal, 1785. Facsimile published in 1972 by Chelsea Publishing Company, New York. (pages 9, 10, and 18)

[61] V. Conitzer. Anonymity-proof voting rules. In *Proceedings of the 4th Conference on Web and Internet Economics (WINE)*, pages 295–306, 2008. (page 218)

[62] V. Conitzer. Should social network structure be taken into account in elections? *Mathematical Social Sciences*, 64(1):100–102, 2012. (pages 19 and 196)

[63] V. Conitzer. The maximum likelihood approach to voting on social networks. In *Proceedings of the 51st Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 1482–1487, 2013. (pages 19, 187, 188, and 196)

[64] V. Conitzer and T. Sandholm. Universal voting protocol tweaks to make manipulation hard. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 781–788, 2003. (page 10)

[65] V. Conitzer and T. Sandholm. Communication complexity of common voting rules. In *Proceedings of the 6th ACM Conference on Economics and Computation (EC)*, pages 78–87, 2005. (pages 18 and 196)

[66] V. Conitzer and T. Sandholm. Common voting rules as maximum likelihood estimators. In *Proceedings of the 21st Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 145–152, 2005. (pages 11, 18, 143, 172, and 174)

[67] V. Conitzer and T. Sandholm. Nonexistence of voting rules that are usually hard to manipulate. In *Proceedings of the 21st AAAI Conference on Artificial Intelligence (AAAI)*, pages 627–634, 2006. (page 10)

[68] V. Conitzer, A. Davenport, and H. Kalagnanam. Improved bounds for computing Kemeny rankings. In *Proceedings of the 21st AAAI Conference on Artificial Intelligence (AAAI)*, pages 620–626, 2006. (page 10)

[69] V. Conitzer, T. Sandholm, and J. Lang. When are elections with few candidates hard to manipulate? *Journal of the ACM*, 54(3):1–33, 2007. (page 10)

[70] V. Conitzer, M. Rognlie, and L. Xia. Preference functions that score rankings and maximum likelihood estimation. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 109–115, 2009. (pages 18, 174, 183,

and 196)

[71] D. E. Critchlow, M. A. Fligner, and J. S. Verducci. Probability models on rankings. *Journal of Mathematical Psychology*, 35(3):294–318, 1991. (page 183)

[72] M. H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974. (pages 19 and 197)

[73] A. Demers, S. Keshav, and S. Shenker. Analysis and simulation of a fair queueing algorithm. In *Proceedings of the ACM Symposium on Communications Architectures & Protocols (SIGCOMM)*, pages 1–12, 1989. (pages 15 and 112)

[74] D. Desario and S. Robins. Generalized solid-angle theory for real polytopes. *The Quarterly Journal of Mathematics*, 62(4):1003–1015, 2011. (page 171)

[75] P. Diaconis and R. L. Graham. Spearman's footrule as a measure of disarray. *Journal of the Royal Statistical Society Series B*, 32(24):262–268, 1977. (page 206)

[76] S. Dobzinski and A. D. Procaccia. Frequent manipulability of elections: The case of two voters. In *Proceedings of the 4th Conference on Web and Internet Economics (WINE)*, pages 653–664, 2008. (page 10)

[77] C. L. Dodgson. *A Method for Taking Votes on More than Two Issues*. Clarendon Press, 1876. (pages 10 and 180)

[78] D. Dolev, D. G. Feitelson, J. Y. Halpern, R. Kupferman, and N. Linial. No justified complaints: On fair sharing of multiple resources. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference (ITCS)*, pages 68–75, 2012. (pages 15, 88, and 112)

[79] C. Dwork, R. Kumar, M. Naor, and D. Sivakumar. Rank aggregation methods for the web. In *Proceedings of the 10th International World Wide Web Conference (WWW)*, pages 613–622, 2001. (pages 10 and 214)

[80] E. Eisenberg and D. Gale. Consensus of subjective probabilities: The pari-mutuel method. *The Annals of Mathematical Statistics*, 30(1):165–168, 1959. (page 25)

[81] E. Elkind and G. Erdélyi. Manipulation under voting rule uncertainty. In *Proceedings of the 11th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 627–634, 2012. (page 160)

[82] E. Elkind and N. Shah. Electing the most probable without eliminating the irrational: Voting over intransitive domains. In *Proceedings of the 30th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 182–191, 2014. (pages 11, 18, and 21)

[83] E. Elkind and A. Slinko. Rationalizations of voting rules. In F. Brandt, V. Conitzer, U. Endress, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 8. Cambridge University Press, 2016. (pages 11, 143, and 183)

[84] E. Elkind, P. Faliszewski, and A. Slinko. On distance rationalizability of some voting rules. In *Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, pages 108–117, 2009. (page 183)

[85] E. Elkind, P. Faliszewski, and A. Slinko. On the role of distances in defining voting rules. In *Proceedings of the 9th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 375–382, 2010. (pages 11, 18, 143, 183, and 184)

[86] E. Elkind, P. Faliszewski, and A. Slinko. Good rationalizations of voting rules. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*, pages 774–779, 2010. (page 196)

[87] U. Endriss and U. Grandi. Binary aggregation by selection of the most representative voters. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, pages 668–674, 2014. (page 183)

[88] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. How hard is bribery in elections? *Journal of Artificial Intelligence Research*, 35:485–532, 2009. (page 10)

[89] "P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. The shield that never was: Societies with single-peaked preferences are more open to manipulation and control. In *Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, pages 118–127, 2009. (page 10)

[90] P. Faliszewski, E. Hemaspaandra, L. A. Hemaspaandra, and J. Rothe. Llull and Copeland voting computationally resist bribery and constructive control. *Journal of Artificial Intelligence Research*, 35:275–341, 2009. (page 10)

[91] P. C. Fishburn. Paradoxes of voting. *American Political Science Review*, 68(2):537–546, 1974. (page 160)

[92] M. A. Fligner and J. S. Verducci. Distance based ranking models. *Journal of the Royal Statistical Society B*, 48(3):359 – 369, 1986. (pages 161 and 183)

[93] D. Foley. Resource allocation and the public sector. *Yale Economics Essays*, 7:45–98, 1967. (pages 6 and 25)

[94] G. Freitas. Combinatorial assignment under dichotomous preferences. Manuscript, 2010. (pages 9, 13, 53, and 55)

[95] R. M. Freund and S. Mizuno. *Interior point methods: current status and future directions*. Springer, 2000. (page 61)

[96] R. M. Freund, R. Roundy, and M. Todd. Identifying the set of always-active constraints in a system of linear inequalities by a single linear program. Technical Report 1674-85, Massachusetts Institute of Technology (MIT), Sloan School of Management, 1985. (page 61)

[97] E. J. Friedman, A. Ghodsi, S. Shenker, and I. Stoica. Strategyproofness, Leontief economies and the Kalai-Smorodinsky solution. Manuscript, 2011. (pages 15, 80, and 87)

[98] D. Gale. *The Theory of Linear Economic Models*. University of Chicago Press, 1960. (pages 8 and 9)

[99] A. Ghodsi, M. Zaharia, B. Hindman, A. Konwinski, S. Shenker, and I. Stoica. Dominant resource fairness: Fair allocation of multiple resource types. In *Proceedings of*

the 8th USENIX Conference on Networked Systems Design and Implementation (NSDI), pages 24–37, 2011. (pages 4, 8, 9, 13, 14, 26, 45, 52, 67, 68, 69, 70, 80, 81, 86, 87, 88, 90, 95, 99, 112, and 113)

[100] A. Ghodsi, V. Sekar, M. Zaharia, and I. Stoica. Multi-resource fair queueing for packet processing. In *Proceedings of the ACM Symposium on Communications Architectures & Protocols (SIGCOMM)*, pages 1–12, 2012. (page 112)

[101] A. Ghodsi, M. Zaharia, S. Shenker, and I. Stoica. Choosy: max-min fair sharing for datacenter jobs with constraints. In *Proceedings of the 8th ACM European Conference on Computer Systems (EuroSys)*, pages 365–378, 2013. (pages 4, 9, 13, 52, and 112)

[102] A. Gibbard. Manipulation of voting schemes. *Econometrica*, 41:587–602, 1973. (page 10)

[103] K. Y. Goldberg, T. Roeder, D. Gupta, and C. Perkins. Eigentaste: A constant time collaborative filtering algorithm. *Information Retrieval*, 4(2):133–151, 2001. (page 133)

[104] J. Goldman and A. D. Procaccia. Spliddit: Unleashing fair division algorithms. *SIGecom Exchanges*, 13(2):41–46, 2014. (pages 26, 43, and 118)

[105] L. Gourvès, J. Monnot, and L. Tlilane. Near fairness in matroids. In *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI)*, pages 393–398, 2014. (page 7)

[106] T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973. (page 6)

[107] V. Guruswami. *List Decoding of Error-Correcting Codes*. Springer, 2005. (pages 20, 200, and 217)

[108] A. Gutman and N. Nisan. Fair allocation without trade. In *Proceedings of the 11th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 719–728, 2012. (pages 15, 88, and 112)

[109] B. Haeupler. Interactive channel capacity revisited. In *Proceedings of the 55th Symposium on Foundations of Computer Science (FOCS)*, pages 226–235, 2014. (page 217)

[110] T. L. Heath. *Diophantus of Alexandria: A study in the history of Greek algebra*. CUP Archive, 2012. (page 205)

[111] J. L. Hellerstein. Google cluster data. Google Research Blog, 2010. Posted at http://googleresearch.blogspot.com/2010/01/google-cluster-data.html. (page 110)

[112] E. Hemaspaandra, L. A. Hemaspaandra, and J. Rothe. Hybrid elections broaden complexity-theoretic resistance to control. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1308–1314, 2007. (page 10)

[113] A. Hylland and R. Zeckhauser. The efficient allocation of individuals to positions. *The Journal of Political Economy*, 87(2):293–314, 1979. (pages 8, 9, 13, and 65)

[114] A. X. Jiang, A. D. Procaccia, Y. Qian, N. Shah, and M. Tambe. Defender (mis)coordination in security games. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 220–226, 2013. (page 21)

[115] A. X. Jiang, L. S. Marcolino, A. D. Procaccia, T. Sandholm, N. Shah, and M. Tambe. Diverse randomized agents vote to win. In *Proceedings of the 28th Annual Conference on Neural Information Processing Systems (NIPS)*, pages 2573–2581, 2014. (pages 21, 193, 194, 195, 247, and 248)

[116] C. Joe-Wong, S. Sen, T. Lan, and M. Chiang. Multi-resource allocation: Fairness-efficiency tradeoffs in a unifying framework. In *Proceedings of the 31st Annual IEEE Conference on Computer Communications (INFOCOM)*, pages 1206–1214, 2012. (page 112)

[117] M. J. Kearns, R. E. Schapire, and L. M. Sellie. Toward efficient agnostic learning. *Machine Learning*, 17:115–141, 1994. (pages 20 and 199)

[118] F. P. Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, 8:33–37, 1997. (pages 12, 41, and 232)

[119] J. G. Kemeny. Mathematics without numbers. *Daedalus*, 88(4):577–591, 1959. (pages 10, 18, and 143)

[120] W. Kets, D. M. Pennock, R. Sethi, and N. Shah. Betting strategies, market selection, and the wisdom of crowds. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, pages 735–741, 2014. (page 21)

[121] L. Khachiyan. A polynomial algorithm in linear programming. *Soviet Mathematics Doklady*, 20:191–194, 1979. (page 62)

[122] J. Kleinberg. Cascading behavior in networks: Algorithmic and economic issues. In N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 24. Cambridge University Press, 2007. (page 197)

[123] F. Kojima. Random assignment of multiple indivisible objects. *Mathematical Social Sciences*, 57(1):134–142, 2009. (pages 13 and 65)

[124] A. Kontorovich. Concentration in unbounded metric spaces and algorithmic stability. In *Proceedings of the 31st International Conference on Machine Learning (ICML)*, pages 28–36, 2014. (page 196)

[125] A. Krause. SFO: A toolbox for submodular function optimization. *Journal of Machine Learning Research*, 11:1141–1144, 2010. (page 141)

[126] D. Kurokawa, A. D. Procaccia, and N. Shah. Leximin allocations in the real world. In *Proceedings of the 16th ACM Conference on Economics and Computation (EC)*, pages 345–362, 2015. (pages 8, 9, 14, 21, and 26)

[127] D. Kurokawa, A. D. Procaccia, and J. Wang. When can the maximin share guarantee be guaranteed? In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, pages 523–529, 2016. (page 26)

[128] S. Lahaie and N. Shah. Neutrality and geometry of mean voting. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*, pages 333–350, 2014. (page 21)

[129] G. Lebanon and J. Lafferty. Cranking: Combining rankings using conditional probability models on permutations. In *Proceedings of the 9th International Confer-*

ence on Machine Learning (ICML), pages 363 – 370, 2002. (page 183)

[130] E. Lee. APX-hardness of maximizing nash social welfare with indivisible items. arXiv:1507.01159, 2015. (pages 9, 12, and 41)

[131] J. Li and J. Xue. Egalitarian division under Leontief preferences. Manuscript, 2011. (pages 15, 80, and 87)

[132] J. Li and J. Xue. Egalitarian division under Leontief preferences. *Economic Theory*, 54(3):597–622, 2013. (page 112)

[133] W. Li, X. Liu, X. Zhang, and X. Zhang. Dynamic fair allocation of multiple resources with bounded number of tasks in cloud computing systems. *Multiagent and Grid Systems*, 11(4):245–257, 2016. (pages 4 and 52)

[134] R. J. Lipton, E. Markakis, E. Mossel, and A. Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 6th ACM Conference on Economics and Computation (EC)*, pages 125–131, 2004. (pages 12, 40, and 82)

[135] T.-Y. Liu. *Learning to Rank for Information Retrieval*. Springer, 2011. (page 183)

[136] T. Lu and C. Boutilier. Robust approximation and incremental elicitation in voting protocols. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 287–293, 2011. (pages 119 and 196)

[137] T. Lu and C. Boutilier. Learning Mallows models with pairwise preferences. In *Proceedings of the 28th International Conference on Machine Learning (ICML)*, pages 145–152, 2011. (pages 148 and 183)

[138] T. Lu and C. Boutilier. Budgeted social choice: From consensus to personalized decision making. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 280–286, 2011. (pages 17, 119, and 142)

[139] C. L. Mallows. Non-null ranking models. *Biometrika*, 44:114–130, 1957. (pages 18, 143, 147, 161, 198, and 217)

[140] A. Mao, A. D. Procaccia, and Y. Chen. Better human computation through principled voting. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI)*, pages 1142–1148, 2013. (pages 11, 18, 20, 143, 144, 161, and 215)

[141] N. Mattei and T. Walsh. Preflib: A library of preference data. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT)*, pages 259–270, 2013. (page 133)

[142] J. C. McCabe-Dansted, G. Pritchard, and A. M. Slinko. Approximability of Dodgson's rule. *Social Choice and Welfare*, 31(2):311–330, 2008. (page 10)

[143] T. Meskanen and H. Nurmi. Closeness counts in social choice. In M. Braham and F. Steffen, editors, *Power, Freedom, and Voting*. Springer-Verlag, 2008. (pages 20, 183, 218, and 239)

[144] B. L. Monroe. Fully proportional representation. *American Political Science Review*, 89(4):925–940, 1995. (page 119)

[145] E. Mossel. Gaussian bounds for noise correlation of functions. *Geometric and Functional Analysis*, 19(6):1713–1756, 2010. (page 197)

[146] E. Mossel, A. D. Procaccia, and M. Z. Rácz. A smooth transition from powerlessness to absolute power. *Journal of Artificial Intelligence Research*, 48:923–951, 2013. (pages 168, 171, 174, 180, and 183)

[147] E. Mossel, J. Neeman, and O. Tamuz. Majority dynamics and aggregation of information in social networks. *Autonomous Agents and Multi-Agent Systems*, 28(3): 408–429, 2014. (pages 19 and 197)

[148] F. Mosteller. Remarks on the method of paired comparisons: I. the least squares solution assuming equal standard deviations and equal correlations. *Psychometrika*, 16(1):3–9, 1951. (page 189)

[149] H. Moulin. *The Strategy of Social Choice*, volume 18 of *Advanced Textbooks in Economics*. North-Holland, 1983. (page 172)

[150] H. Moulin. Uniform externalities: Two axioms for fair allocation. *Journal of Public Economics*, 43(3):305–326, 1990. (page 7)

[151] H. Moulin. *Fair Division and Collective Welfare*. MIT Press, 2003. (pages 5 and 26)

[152] H. Moulin and R. Stong. Fair queuing and other probabilistic allocation methods. *Mathematics of Operations Research*, 27(1):1–30, 2002. (pages 15, 82, and 112)

[153] D. Nace and J. B. Orlin. Lexicographically minimum and maximum load linear programming problems. *Operations research*, 55(1):182–187, 2007. (page 60)

[154] J. Nash. The bargaining problem. *Econometrica*, 18(2):155–162, 1950. (pages 12 and 41)

[155] T. T. Nguyen, M. Roos, and J. Rothe. A survey of approximability and inapproximability results for social welfare optimization in multiagent resource allocation. *Annals of Mathematics and Artificial Intelligence*, 68(1):65–90, 2013. (pages 12, 37, 38, and 234)

[156] A. Nicolò. Efficiency and truthfulness with Leontief preferences. A note on two-agent, two-good economies. *Review of Economic Design*, 8(4):373–382, 2004. (page 68)

[157] A. Nongaillard, P. Mathieu, and B. Jaumard. A realistic approach to solve the Nash welfare. In *Proceedings of the 7th International Conference on Practical Applications of Agents and Multi-Agent Systems (PAAMS)*, pages 374–382, 2009. (page 37)

[158] S. Obraztsova, E. Elkind, and N. Hazon. Ties matter: Complexity of voting manipulation revisited. In *Proceedings of the 10th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 71–78, 2011. (page 155)

[159] S. Obraztsova, E. Elkind, P. Faliszewski, and A. Slinko. On swap-distance geometry of voting rules. In *Proceedings of the 12th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 383–390, 2013. (page 183)

[160] A. Othman, T. Sandholm, and E. Budish. Finding approximate competitive equilibria: Efficient and fair course allocation. In *Proceedings of the 9th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 873–880, 2010. (page 86)

[161] C. H. Papadimitriou and M. Yannakakis. Optimization, approximation, and complexity classes. *Journal of Computer and System Sciences*, 43(3):425–440, 1991. (page 135)

[162] V. Pareto. *Manuale di economia politica: Con una introduzione alla scienza sociale*. CEDAM, 1974. (page 5)

[163] D. C. Parkes. Online mechanisms. In N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 16. Cambridge University Press, 2007. (page 88)

[164] D. C. Parkes, A. D. Procaccia, and N. Shah. Beyond Dominant Resource Fairness: Extensions, limitations, and indivisibilities. In *Proceedings of the 13th ACM Conference on Economics and Computation (EC)*, pages 808–825, 2012. (page 72)

[165] D. C. Parkes, A. D. Procaccia, and N. Shah. Beyond Dominant Resource Fairness: Extensions, limitations, and indivisibilities. *ACM Transactions on Economics and Computation*, 3(1): article 3, 2015. (pages 4, 7, 9, 13, 21, 45, 52, 90, 112, and 113)

[166] E. Pazner and D. Schmeidler. Egalitarian equivalent allocations: A new concept of economic equity. *Quarterly Journal of Economics*, 92(4):671–687, 1978. (pages 41 and 46)

[167] T. Pfeiffer, X. A. Gao, A. Mao, Y. Chen, and D. G. Rand. Adaptive polling for information aggregation. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI)*, pages 122–128, 2012. (page 183)

[168] A. D. Procaccia. Can approximation circumvent Gibbard-Satterthwaite? In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*, pages 836–841, 2010. (page 218)

[169] A. D. Procaccia. Cake cutting: Not just child's play. *Communications of the ACM*, 56(7):78–87, 2013. (pages 8 and 45)

[170] A. D. Procaccia. Cake cutting algorithms. In F. Brandt, V. Conitzer, U. Endress, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 13. Cambridge University Press, 2016. (page 8)

[171] A. D. Procaccia and J. S. Rosenschein. The distortion of cardinal preferences in voting. In *Proceedings of the 10th International Workshop on Cooperative Information Agents (CIA)*, pages 317–331, 2006. (pages 11, 17, 117, 119, and 142)

[172] A. D. Procaccia and N. Shah. Is approval voting optimal given approval votes? In *Proceedings of the 29th Annual Conference on Neural Information Processing Systems (NIPS)*, pages 1792–1800, 2015. (pages 11, 18, and 21)

[173] A. D. Procaccia and N. Shah. Optimal aggregation of uncertain preferences. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, pages 608–

614, 2016. (page 21)

[174] A. D. Procaccia and M. Tennenholtz. Approximate mechanism design without money. In *Proceedings of the 10th ACM Conference on Economics and Computation (EC)*, pages 177–186, 2009. (page 78)

[175] A. D. Procaccia and J. Wang. Fair enough: Guaranteeing approximate maximin shares. In *Proceedings of the 14th ACM Conference on Economics and Computation (EC)*, pages 675–692, 2014. (pages 8, 26, and 31)

[176] A. D. Procaccia, J. S. Rosenschein, and G. A. Kaminka. On the robustness of preference aggregation in noisy environments. In *Proceedings of the 6th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 416–422, 2007. (page 218)

[177] A. D. Procaccia, J. S. Rosenschein, and A. Zohar. Multi-winner elections: Complexity of manipulation, control and winner-determination. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1476–1481, 2007. (page 10)

[178] A. D. Procaccia, J. S. Rosenschein, and A. Zohar. On the complexity of achieveing proportional representation. *Social Choice and Welfare*, 30(3):353–362, 2008. (pages 17, 119, and 142)

[179] A. D. Procaccia, S. J. Reddi, and N. Shah. A maximum likelihood approach for selecting sets of alternatives. In *Proceedings of the 28th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 695–704, 2012. (pages 4, 11, 18, 21, 118, 143, 184, and 196)

[180] A. D. Procaccia, N. Shah, and M. L. Tucker. On the structure of synergies in cooperative games. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, pages 763–769, 2014. (page 21)

[181] A. D. Procaccia, N. Shah, and E. Sodomka. Ranked voting on social networks. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2040–2046, 2015. (page 21)

[182] A. D. Procaccia, N. Shah, and Y. Zick. Voting rules as error-correcting codes. *Artificial Intelligence*, 231:1–16, 2016. Preliminary version in AAAI'15. (pages 11 and 21)

[183] M. Pycia. Assignment with multiple-unit demands and responsive preferences. Manuscript, 2011. (pages 13 and 65)

[184] S. Ramezani and U. Endriss. Nash social welfare in multiagent resource allocation. In *Proceedings of the 12th International Workshop on Agent-Mediated Electronic Commerce (AMEC)*, pages 117–131, 2010. (pages 12, 26, 37, and 40)

[185] B. Rochwerger, D. Breitgand, E. Levy, A. Galis, K. Nagin, I. Llorente, R. Montero, Y. Wolfsthal, E. Elmroth, J. Cáceres, M. Ben-Yehuda, W. Emmerich, and F. Galán. The RESERVOIR model and architecture for open federated cloud computing. *IBM Journal of Research and Development*, 53(4), 2009. (page 67)

[186] A. E. Roth, T. Sönmez, and M. U. Ünver. Pairwise kidney exchange. *Journal of Economic Theory*, 125:151–188, 2005. (pages 4, 9, 13, and 52)

[187] D. G. Saari. *Basic geometry of voting*. Springer, 1995. (page 183)

[188] D. G. Saari. Complexity and the geometry of voting. *Mathematical and Computer Modelling*, 48(910):1335–1356, 2008. (page 183)

[189] M. Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10:187–217, 1975. (page 10)

[190] J. Schummer. Strategy-proofness versus efficiency on restricted domains of exchange economies. *Social Choice and Welfare*, 14:47–56, 1997. (page 42)

[191] P. Skowron, P. Faliszewski, and J. Lang. Finding a collective set of items: From proportional multirepresentation to group recommendation. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, pages 2131–2137, 2015. (pages 17 and 142)

[192] Y. Tahir, S. Yang, A. Koliousis, and J. McCann. UDRF: Multi-resource fairness for complex jobs with placement constraints. In *Proceedings of the IEEE Global Communications Conference (GLOBECOM)*, pages 1–7, 2015. (pages 4, 13, and 52)

[193] E. Tardos. A strongly polynomial algorithm to solve combinatorial linear programs. *Operations Research*, 34(2):250–256, 1986. (page 61)

[194] L. L. Thurstone. A law of comparative judgement. *Psychological Review*, 34:273–286, 1927. (page 189)

[195] H. Varian. Equity, envy and efficiency. *Journal of Economic Theory*, 9:63–91, 1974. (pages 3, 8, 9, 25, 28, and 82)

[196] W. Vickrey. Counter speculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961. (page 6)

[197] J. von Neumann. A certain zero-sum two-person game equivalent to the optimal assignment problem. In W. Kuhn and A. W. Tucker, editors, *Contributions to the Theory of Games*, volume 2, pages 5–12. Princeton University Press, 1953. (pages 59, 63, 65, 126, and 131)

[198] T. Walsh. Online cake cutting. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT)*, pages 292–305, 2011. (pages 15, 94, and 111)

[199] W. Wang, B. Li, and B. Liang. Dominant resource fairness in cloud computing systems with heterogeneous servers. In *Proceedings of the 33rd Annual IEEE Conference on Computer Communications (INFOCOM)*, pages 583–591, 2014. (pages 4, 13, and 52)

[200] W. Wang, B. Li, B. Liang, and J. Li. Towards multi-resource fair allocation with placement constraints. In *Proceedings of the 2016 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, pages 415–416, 2016. (pages 4, 13, and 52)

[201] D. Weller. Fair division of a measurable space. *Journal of Mathematical Economics*, 14(1):5–17, 1985. (pages 3, 8, and 9)

[202] E. P. Wigner. The unreasonable effectiveness of mathematics in the natural sciences. *Communications on Pure and Applied Mathematics*, 13(1):1–14, 1960. (page 26)

[203] L. Xia. Generalized scoring rules: a framework that reconciles Borda and Condorcet. *SIGecom Exchanges*, 12(1):42–48, 2013. (page 145)

[204] L. Xia and V. Conitzer. Generalized scoring rules and the frequency of coalitional manipulability. In *Proceedings of the 9th ACM Conference on Economics and Computation (EC)*, pages 109–118, 2008. (pages 145, 174, 180, and 184)

[205] L. Xia and V. Conitzer. Finite local consistency characterizes generalized scoring rules. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 336–341, 2009. (pages 145 and 183)

[206] L. Xia and V. Conitzer. A maximum likelihood approach towards aggregating partial orders. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 446–451, 2011. (pages 18, 184, and 196)

[207] L. Xia, V. Conitzer, and J. Lang. Aggregating preferences in multi-issue domains by using maximum likelihood estimators. In *Proceedings of the 9th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 399–408, 2010. (pages 18 and 196)

[208] H. P. Young. Social choice scoring functions. *SIAM Journal of Applied Mathematics*, 28(4):824–838, 1975. (page 183)

[209] H. P. Young. Condorcet's theory of voting. *The American Political Science Review*, 82(4):1231–1244, 1988. (pages 4, 10, 18, 118, 143, and 196)

[210] H. P. Young and A. Levenglick. A consistent extension of Condorcet's election principle. *SIAM Journal on Applied Mathematics*, 35(2):285–300, 1978. (pages 172 and 174)

[211] S. M. Zahedi and B. C. Lee. REF: Resource elasticity fairness with sharing incentives for multiprocessors. In *Proceedings of the 19th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pages 145–160, 2014. (page 112)

[212] J. Zou, S. Gujar, and D. C. Parkes. Tolerable manipulability in dynamic assignment without money. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*, pages 947–952, 2010. (page 88)

[213] W. Zwicker. Introduction to the theory of voting. In F. Brandt, V. Conitzer, U. Endress, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 2. Cambridge University Press, 2016. (page 10)