# Fair and Efficient Social Decision Making

by

**Paraskevi (Evi) Micha**

A thesis submitted in conformity with the requirements
for the degree of Doctor of Philosophy
Graduate Department of Computer Science
University of Toronto

© Copyright by Paraskevi (Evi) Micha (2023)

## Abstract

Fair and Efficient Social Decision Making

Paraskevi (Evi) Micha

Doctor of Philosophy, Graduate Department of Computer Science, University of Toronto, 2023

Designing decision-making algorithms that *treat people fairly*, *use limited resources efficiently*, and *foster social good* is of high importance. We contribute toward this direction by optimizing decision-making for social good, analyzing and recommending improvements to methods of civic participation, and suggesting novel notions of fairness for domains like machine learning and peer review.

In Part I, we focus on optimizing decision-making for social good in two real-life applications. First, we consider the problem of *allocating pooled diagnostic tests* during an epidemic. Second, we analyze the effect of different grading systems on student performance and identify conditions under which certain grading systems are optimal.

In Part II, we study three approaches used to increase civic participation. In the first approach, called *sortition*, we suggest selection algorithms that satisfy fairness and efficiency guarantees. In the second approach, called *participatory budgeting*, we formalize the idea of individual fairness in participatory budgeting and study to what extent it can be guaranteed. The last approach is to conduct *polls* for an election to better engage citizens, and we study the effectiveness of polls for predicting the outcome of ranked-choice elections.

In Part III, we leverage the economic literature to propose meaningful notions of fairness for novel application domains. First, inspired by the idea of *core* in economics, we provide approximate fairness guarantees for *clustering*. Second, we utilize *Nash welfare* to design fair algorithms for a *multi-agent variant of the classical multi-armed bandits problem*. Third, we adopt a fairness notion from the fair division literature, namely *envy-freeness*, to *two-sided matching markets*. Fourth, we also import core, to the problem of *reviewer assignment* and to the problem of *team formation*.

**Acknowledgements**

First and foremost, I would like to thank Nisarg Shah, who for the past five years has been the best advisor I could have asked for. Thank you, Nisarg, for teaching me how to approach technical questions, select elegant problems, explore new directions, and write papers effectively, among so many other things. Your passion and enthusiasm for research during every meeting was a constant source of motivation, inspiring me to continuously strive towards becoming a better researcher. I am forever grateful to you for being there when I needed you (even under difficult circumstances), for your guidance throughout these years, and for your friendship. I hope to be a mentor who can offer even a fraction of the support that you have provided me.

To my regular thesis committee, Allan Borodin and Toniann Pitassi, thank you for the valuable feedback and guidance you have provided me over the years, helping me make informed decisions about my career path. To my external examiner, Edith Elkind, thank you for your comments that have greatly improved this thesis. Lastly, I would like to thank Peter Marbach for serving as a committee member.

It goes without saying that this thesis would not exist without the invaluable contributions of the wonderful people with whom I have had the privilege of collaborating: Haris Aziz, Siddhartha Banerjee, Markus Brill, Ioannis Caragiannis, Soroush Ebadian, Rupert Freeman, Aris Filos-Ratsikas, Simon Finster, Aditya Ganguly, Vasilis Gkatzelis, Michelle González Amador, Safwan Hossain, Jonas Israel, Billy Jin, Lily Li, Edwin Lock, Francisco Marmolejo-Cossío, Aleksandar Nikolov, Jannik Peters, Ariel D. Procaccia, Shreyas Sekar, Nisarg Shah, and Alexandros Voudouris.

Special thanks to Ioannis Caragiannis for introducing me to the fascinating world of research, for directing me towards graduate studies, and for his unwavering support and guidance throughout the years. I also want to thank Haris Aziz and Ariel D. Procaccia for hosting me during two fertile and fun semesters in their labs.

I would like to extend a heartfelt thank you to all my friends who made my time in Toronto truly special. I was fortunate to have met many of them from the very beginning of my journey, which greatly eased my transition in a country so far away from home. Last but certainly not least, I am grateful to my family for always supporting and encouraging me to pursue my dreams.

# Contents

**1 Introduction** **1**
1.1 Thesis Structure and Our Results . . . 3
1.2 Bibliographic Notes . . . 12

## I Optimizing Decision-Making for Social Good 14

**2 Welfare-Maximizing Pooled Testing** **15**
2.1 Introduction . . . 15
2.2 Model . . . 17
2.3 Theoretical Results . . . 18
2.4 Practical Algorithmic Implementations . . . 33
2.5 Empirical Results . . . 36
2.6 Discussion . . . 37

**3 Grading Schemes and Students' Performance in Consecutive Tests** **41**
3.1 Introduction . . . 41
3.2 Model . . . 44
3.3 Uniform Letter Grading vs Numerical Scoring . . . 45
3.4 Experiments . . . 59
3.5 Discussion . . . 61

## II Civic Participation 67

**4 Fairness and Representation in Sortition** **68**
4.1 Introduction . . . 68
4.2 Model . . . 69
4.3 Representation Through $q$-Social Cost . . . 70
4.3.1 Representation with Perfect Fairness for $q > k/2$ . . . 70
4.3.2 Representation with Relaxed Fairness for $q \leqslant k/2$ . . . 72
4.4 Representation Through the $q$-Core . . . 75
4.4.1 Fairness and The Ex Post Core . . . 77
4.4.2 Uniform Selection and the Ex Ante Core . . . 82

4.5 Discussion . . . 84

**5 Individual Fairness in Participatory Budgeting 86**
5.1 Introduction . . . 86
5.2 Preliminaries . . . 87
5.3 Individual Representation (IR) . . . 88
5.3.1 ABC Rules Violating IR . . . 91
5.3.2 Computational Complexity . . . 92
5.4 Domain Restrictions . . . 93
5.5 Experimental Results . . . 100
5.6 Discussion . . . 101

**6 Effectiveness of Polls for Predicting Ranked Election Outcomes 102**
6.1 Introduction . . . 102
6.2 Preliminaries . . . 104
6.3 Worst-Case Predictability . . . 105
6.4 Average-Case Predictability . . . 111
6.5 Experiments . . . 117
6.6 Winner Prediction . . . 118
6.7 Discussion . . . 121

**III Algorithmic Fairness 122**

**7 Proportionally Fair Clustering 123**
7.1 Introduction . . . 123
7.2 Preliminaries . . . 125
7.3 Greedy Capture . . . 125
7.4 Universal Lower Bounds . . . 128
7.5 Clustering in Graphs . . . 130
7.6 Computational Aspects . . . 132
7.7 Learning Fair Clustering . . . 139
7.8 Discussion . . . 140

**8 Fair Algorithms for Multi-Agent Multi-Armed Bandits 142**
8.1 Introduction . . . 142
8.2 Preliminaries . . . 146
8.3 The Lower Bound . . . 147
8.4 Explore-First . . . 148
8.5 Epsilon-Greedy . . . 152
8.6 UCB . . . 158
8.7 Discussion . . . 161

**9 Two-sided Markets Meet Fair Division 163**

9.1 Introduction . . . 163
9.2 Preliminaries . . . 165
9.3 Double Envy-Freeness Up To One Match . . . 167
9.4 Double Maximin Share Guarantee . . . 177
9.5 Simulations . . . 180
9.6 Discussion . . . 181

**10 Group Fairness in Peer Review** **183**
10.1 Introduction . . . 183
10.2 Model . . . 186
10.3 CoBRA: An Algorithm for Computing Core-Based Reviewer Assignment . . . 188
10.3.1 Description of CoBRA . . . 191
10.3.2 Main Result . . . 192
10.4 The Core and Other Objective . . . 196
10.5 Experiments . . . 200
10.6 Discussion . . . 202

**11 Fair Team Formation** **204**
11.1 Introduction . . . 204
11.2 Preliminaries . . . 206
11.3 Core . . . 208
11.4 Envy-Freeness . . . 214
11.5 Beyond Balancedness . . . 215
11.5.1 Core . . . 215
11.5.2 Envy-Freeness . . . 217
11.6 Discussion . . . 218

**IV Conclusion** **219**

# Chapter 1

# Introduction

Algorithms and artificial intelligence (AI) are increasingly affecting human lives as they are extensively used for automating critical decisions. To give a few examples, they are often used to decide if an email is spam, who may be a good mate for someone, who will give someone a ride to the airport, or even which book matches someone's taste. In many cases, algorithmic decisions have the potential to impact not only individual users but also groups of people. For example, the decision of who gets the only cab that is nearby a particular neighborhood may affect two or more people, while the decisions of pricing algorithms affect all the potential customers. The common ground in many of these applications is that different people or users have different needs or preferences that an algorithm needs to take into account for making a "good" algorithmic decision.

The concept of a good decision can vary depending on the context. In some cases, the goal might be to optimize accuracy, such as in the design of a spam filtering algorithm, while in other cases efficiency might be of higher importance, such as in the design of book matching algorithms. Fairness and equity are two other factors that can play a huge rule in different applications, such as in the design of riding algorithms. In other cases a balance among accuracy, efficiency, fairness, equity, and privacy might be necessary. As algorithmic decision-making systems become increasingly integrated into society, there is a growing awareness of the critical need to design algorithms that not only prioritize these properties but can also be provably demonstrated to satisfy them. In this thesis, we contribute toward this direction by focusing on

> *... the design of decision-making algorithms that are provably fair (at the level of individuals and groups) and efficient (in using limited resources), and foster social good.*

The work that is presented in this thesis focuses on applications where algorithmic decisions affect groups of people that have different preferences or needs. The aggregation of heterogeneous preferences or needs of a set of people towards a collective outcome has been extensively studied in social choice theory, a traditional field of microeconomics. However, the modern field of computational social choice incorporates the computational perspective on this problem using tools from theoretical computer science. The work in this thesis builds on the field of computational social choice by taking two different approaches:

- first, by considering problems that traditionally fall into the domain of computational social choice, analyzing methods and designing algorithms to effectively tackle them,

- and second, by applying ideas that emerge from computational social choice to different domains like machine learning.

Below, we provide a brief overview of the field of computational social choice and highlight the contribution of this thesis on both categories.

## Computational Social Choice and High Level Contribution of Thesis

Computational social choice is an interdisciplinary field that combines tools from social choice theory and theoretical computer science. Similarly with social choice theory, computational social choice aims to address the problem of aggregating heterogeneous preferences to collective decisions, but instead of focusing mainly on an axiomatic approach, it also utilizes tools from theoretical computer science like complexity theory, approximation algorithms, and worst-case analysis [Brandt et al., 2016].

One central topic in computational social choice is fair division and resource allocation where a scarce resource must be allocated to a set of agents with different needs or preferences. Applications of this topic include the allocation of limited resources during a crisis such as a pandemic. Another central topic is voting theory, which studies the design and analysis of voting rules that take as input the heterogeneous opinions, such as rankings over alternatives, of voters/agents, and output a final decision, such as an aggregate ranking or a winner. Modern applications of this topic include democratic innovations that aim to increase civic participation in our current democratic system. An example of such democratic innovations is participatory budgeting where citizens express their opinions on how a public budget should be spent, leading to collective decision-making. Another example is sortition, where citizens' assemblies– panels randomly selected– are in charge of deliberating and making recommendations about important topics.

In economics and social choice theory, and as an extension in computational social choice, a central theme is the concept of *fairness*. Various fairness criteria have been proposed. For example, the fair division literature offers compelling fairness notions such as proportionality [Steinhaus, 1948], envy-freeness [Foley, 1967], equitability and maxmin fair share [Budish, 2011] that have been studied since the 1950s. Another notion of fairness that derives from cooperative game theory is the core (sometimes known as core stability) [Varian, 1974]. Two other notions of fairness that derive from economics are Nash social welfare maximization [Kaneko and Nakamura, 1979] and Rawlsian egalitarian fairness [Rawls, 1971].

As mentioned above, in this thesis, we build on the field of computational social choice in two different ways. On the one hand, we consider domains such as resource allocation and voting, to which algorithms emerging from computational social choice are traditionally applied. We propose novel frameworks for viewing these traditional domains in a new light, and use these frameworks to design improved algorithms for their real-world applications. On the other hand, we also explore new domains of algorithmic decision-making where notions of fairness that derive from computational social choice are surprisingly effective, such as machine learning applications. Below, we briefly describe the contribution of this thesis in both categories.

In the first category, we focus on the question of how we should make collective decisions, given the heterogeneous preferences, opinions, or needs of groups of people and limited resources. We further distinguish into two distinct categories. First, we pay special attention to real-life applications where the goal is to make optimized decisions that foster the social good by maximizing the social welfare.

More specifically, we focus on two specific applications: one related to allocating pooled diagnostic tests during an epidemic, and another related to examining how different grading systems can impact students' performance. Second, we turn our attention to democratic innovations that foster civic participation. In this thesis, we particularly focus on sortition, participatory budgeting and polls. Our goals include (a) defining desirable properties that procedures for improving civic participation should possess, (b) evaluating whether existing procedures meet these criteria, and (c) suggesting potential improvements.

In the second category, we focus on leveraging the literature on economics and social choice, and proposing meaningful notions of fairness in various domains such as machine learning, matching, peer review, and more. In many cases, traditional machine learning algorithms make decisions that affect a whole population. To give some examples, machine learning algorithms are used to design a non-personalized recommender system or a clustering solution that categorizes users to allocate some resources accordingly. Such applications have induced the need to design *fair* machine learning algorithms, and this has been a subject of intense study in the last decade [Mehrabi et al., 2021]. However, while fairness is a nascent subject within the machine learning literature, it has a long history in computational social choice. In this thesis, we aim to show how notions of fairness from computational social choice can effectively be applied to classical machine learning problems, such as clustering, reinforcement learning, and recommender systems, as well as other important applications such as peer review assignment and team formation. Our goals are (a) to suggest meaningful and non-trivial notions of fairness for the application domains under consideration that are inspired from economics and (b) investigate when and how they can be (approximately) guaranteed.

## 1.1 Thesis Structure and Our Results

The content of this thesis is organized in three parts. Part I describes the work on optimizing decision-making for social good, Part II describes the work on civic participation, and Part III describes the work on algorithmic fairness.

### Part I: Optimizing Decision-Making For Social Good

The first part of the thesis focuses on how algorithms can be utilized to make decisions for social good. In recent years, AI algorithms have played a significant role in addressing complex societal problems related to public health, education, conservation, disaster management, homelessness, and more. Making decisions that foster social good is of high importance in most of these applications. But what does social good mean? At a high level, social good refers to actions that aim to benefit individuals, communities, or society as a whole, but the specific definition is shaped by the unique requirements of each application. Concepts associated with social good include social welfare, individual and group fairness, justice, transparency, inclusiveness, and privacy.

In this part, we consider two real life optimization problems. The first one was significantly motivated and informed by a collaboration with the Potosinian Institute of Scientific and Technological Research (IPICYT), a higher education research institution in Mexico. This problem emerged during the COVID-19 era when all the employees of the institute were in lockdown. The goal was the employees of the institute to return back to in person activities with safety, by allocating the scarce resource of pooled tests. The second problem is an optimization problem related to the choice of the correct grading scheme for

boosting students performance in consecutive examinations. For both problems, the goal is to optimize decisions for social good, where social good is defined through social welfare.

### Chapter 2: Welfare-Maximizing Pooled Testing

This chapter focuses on a real life problem that emerged from our collaboration with IPICYT. In particular, back in November 2021, all the employees of IPICYT were in full lockdown due to COVID-19, and only employees that have been tested negative could have access to the institute for a few days. During the COVID-19 era, the idea of pooled testing was extensively used to increase the speed and cost of testing, a technique that was originally proposed by Dorfman [1943] during the Second World War for screening the US army for syphilis. In a pooled test, multiple samples are pooled together to one test; if the test detects the virus, this means that at least one individual that was pooled in the test is infected, otherwise, all the individuals are healthy. IPICYT had available a limited number of pooled tests every week that it could use for testing its population and allow the employees that were included in a negative test to return back to in person activities.

Inspired by this problem, we asked the following question:

> *In an epidemic, how should an organization with limited testing resources safely return to in-person activities after a period of lockdown?*

We study this question in a setting where the population at hand is heterogeneous in both utility for in-person activities and probability of infection. Under the assumption that samples can be pooled, the question of how to allocate a limited testing budget in the population to maximize the aggregate utility (i.e. welfare) of negatively-tested individuals who return to in-person activities is non-trivial, with a large space of potential testing allocations.

One important observation is that by allowing the sample of an individual to be included into more than one test, we can achieve higher welfare. To see that, consider the case where some individuals are healthy almost with certainty. If we include these individuals to all the tests, the probability of a test to turn negative is not affected and the probability of each of them to be included in a negative test only increases. However, this kind of solutions are logistically more difficult for laboratory personnel. So a reasonable question is: How much do you gain if we allow each sample to be included into multiple tests? We show that this gain is a small constant! In the light of this result, we design an efficient algorithm for finding testing allocations, where each sample is included in at most one test, with approximately optimal welfare.

We tested our algorithm on real data that were gathered from a pilot study executed at IPICYT, in September 2022. In the computational experiments, we highlight the efficacy and viability of our algorithm in practice.

### Chapter 3: Grading Schemes and Students' Performance in Consecutive Tests

In this chapter, we study letter grading schemes, which are routinely employed for evaluating student performance. Typically, a numerical score obtained through one or more evaluations is converted into a letter grade (e.g., A+, B-, etc.) by associating a disjoint interval of numerical scores to each letter grade.

We propose the first model for studying the (de)motivational effects of such grading on the students and, consequently, on their performance in future evaluations. In particular, we assume that each

student has an intrinsic ability, of which she is aware, and her score in an evaluation derives from this intrinsic ability by adding some random noise which captures random factors that may affect student's performance. The student learns the letter grade of the interval that her actual numerical score falls in and perceives her score as the midpoint of this interval, as it happens in practise (see, e.g., [University of Western Ontario, 2022, Victoria University of Wellington, 2022]). Then, if student's score falls below her intrinsic ability, the student gets demotivated, otherwise she gets motivated, as it has been observed [Deci et al., 1999, Dev, 1997, Cameron and Pierce, 1994]. This (de)motivation effect is reflected on the intrinsic quality of the student which changes accordingly. In the next evaluation, student's performance derives form the affected intrinsic quality.

We use this model to compare uniform letter grading schemes, in which the range of scores is divided into equal-length parts that are mapped to the letter grades, to numerical scoring, in which the score is not converted to any letter grade (equivalently, every score is its own letter grade).

Theoretically, we compare the different schemes in the case where two sequential evaluations take place. We show that under natural conditions, when the motivational and demotivational effects are equal, all these grading schemes have equal performance, but otherwise either numerical scoring outperforms all uniform letter grading schemes or the opposite happens. To be able to give a more concrete distinction of the different grading schemes we make further assumptions with the strongest one being that students' intrinsic qualities follow a uniform distribution. Under these assumptions, we can conclude that numerical scoring outperforms all uniform letter grading schemes when the demotivational effect is stronger than the motivational effect, and the opposite happens when the demotivational effect is weaker than the motivational effect.

We also empirically compare numerical scoring to uniform letter grading schemes, by assuming that students' intrinsic qualities derive from a normal distribution. Under two sequential evaluations, we observe the same phenomena as we prove in the theoretical part. However, surprisingly, when more evaluations take place, at some point the effect is reversed, i.e. numerical scoring outperforms all uniform letter grading schemes when the motivational effect is stronger than the demotivational effect, and the opposite happens when the motivational effect is weaker than the demotivational effect. Our experimental results indicate that the choice of the grading scheme depends on the number of consecutive evaluations.

## Part II: Civic Participation

In this part, we ask how algorithms can improve the functioning of democracy. In particular, we focus on three different democratic innovations that are used in practice for increasing civic participation in decision making. The first approach, known as sortition, builds on the idea of creating panels of representatives chosen at random from a given population. These panels are known as *citizens' assemblies* and the participants are in charge of deliberating and making recommendations for a particular topic at hand. The second approach, known as participatory budgeting, gathers and aggregates votes of citizens into collective decisions, when a limited public budget can be allocated. A third approach of incorporating citizens' opinions on policy making is the use of polls. Other democratic innovations that are used in practise, but are out of the scope of this thesis, include liquid democracy and online deliberation platforms. While all these democratic innovations are powerful and have the potential to increase civic participation, their effectiveness heavily relies on the design and implementation of processes that guarantee the fulfillment of desirable properties.

In this part, for each of these three approaches, we mathematically analyze processes that are used in practice to see whether they provide meaningful theoretical guarantees, and in some cases we suggest improved processes that provably satisfy desirable properties. We believe that technology has the power to change the way that we execute democracy and this part contributes toward this direction.

**Chapter 4: Fairness and Representation in Sortition**

*Sortition* is a form of democracy built on random selection of representatives. This idea was used in Athenian democracy in the 6th century BC, as a way of achieving the principal idea of this democracy which was that equals should have equal participation [Van Reybrouck, 2016]. During the American and French revolutions, democracy became synonymous with the election and sortition was forgotten. However, the last decades, the idea of sortition has come back to light and is used around the world, mainly in the form of citizens' assemblies — randomly selected groups of people who deliberate on central questions and make recommendations.

Two of the key arguments in favor of sortition are that it provides *representation* (a random panel reflects the composition of the population) and *fairness* (everyone has a chance to participate). A natural way of applying sortition is to simply apply *uniform selection*, i.e. if the goal is to create a representative panel of size $k$ from a given population of size $n$, select $k$ individuals uniformly at random [Engelstad, 1989]. The main questions that we aim to answer in this chapter are:

- *What is a formal definition of representation of a population?*
- *To what extent does uniform selection satisfy representation?*
- *Is it possible to design selection procedures that enhance representation guarantees while maintaining fairness?*

Towards answering these questions, we introduce the notion of a representation metric on the space of individuals, and assume that the cost of an individual for a panel is determined by the $q$-th closest representative, where $q \in \{1, \ldots, k\}$. To measure the representation of a panel in a population level, we take two different approaches.

The first approach builds on the idea of using the social cost of a panel as a measurement of representation. Formally, the representation of a (random) panel is measured by the ratio between the (expected) sum of costs of the optimal panel for the individuals and that of the given panel. Moving to the second question of whether uniform selection is also representative, our answer is that "it depends" — on the value of $q$. For $k/2 < q \leqslant k - \Omega(k)$, where $k$ is the panel size, we show that uniform random selection is indeed representative by establishing a constant lower bound on this ratio. By contrast, for $q \leqslant k/2$, no random selection algorithm that is almost fair can give such a guarantee. Turning our attention to the last question, we consider relaxed fairness guarantees and develop a new random selection algorithm that sheds light on the tradeoff between representation and fairness.

The second approach builds on the idea of proportionality and specifically on the observation that if a group of individuals contains $x\%$ of the population, then $x\%$ of the panel should consist of individuals from this group. We define this intuitive property formally by using a notion called *core*. At a high level, a panel is in the core if no group of individuals is underrepresented proportional to its size. We show that uniformly random selection of a decision panel satisfies (approximately) *ex ante core*. In

practice, however, it is often required core to be satisfied with certainty. Moving to the last question, as stated above, we ask if we can design a selection algorithm that satisfies fairness and *ex post core* simultaneously. We answer this question affirmatively and present an efficient selection algorithm that is fair and is in the approximately ex post core.

### Chapter 5: Individual Fairness in Participatory Budgeting

Participatory budgeting is a democratic approach that is used by many cities around the world for deciding how to spend a public budget. In particular, citizens vote on how the budget should be allocated to different public projects, and then the local authority takes into account these votes for deciding which projects will be funded. A reasonable question is:

> *How can we fairly aggregate citizens' votes towards a collective decision of allocating a public budget?*

This question erased a body of research in the computational social choice literature related to modelling different participatory budgeting cases, suggesting the design of the ballot preferences and their subsequent aggregation [Aziz et al., 2017]. In this chapter, we focus on the case where the ballots have the form of approval preferences, all projects have the same cost and each of them should either fully be implemented or not be implemented, and the budget is enough for funding up to $k$ projects. This case is well-known as an approval-based committee elections problem [Aziz et al., 2017].

When selecting multiple candidates based on approval preferences of voters, the proportional representation of voters' opinions is an important and well-studied desideratum. Existing criteria for evaluating the representativeness of outcomes focus on *groups* of voters and demand that sufficiently large and cohesive groups are "represented" in the sense that candidates approved by *some* group members are selected. Crucially, these criteria say nothing about the representation of *individual* voters, even if these voters are members of groups that deserve representation.

In this chapter, we formalize the concept of *individual representation (IR)* and explore to which extent, and under which circumstances, it can be achieved. Our concept aims to ensure all members of a voter group equal guarantees. Intuitively, when a population consists of $n$ voters and a committee of $k$ representatives is elected, if a voter is a member of a group of size at least equal to $\ell \cdot n/k$ such that all the members of the group approve at least the same $\ell$ candidates, then the voter can claim at least $\ell$ representatives that she approves to be included in the committee.

We show that, without further assumptions, checking whether an IR outcome exists is computationally intractable, and we verify that all common approval-based voting rules may fail to provide IR even in cases where this is possible. We then focus on domain restrictions [Elkind and Lackner, 2015, Yang, 2019] and show that positive results can be obtained in some cases. Two famous restriction concepts are "voter interval" and "candidate interval" preferences. Interestingly, we show that while for the former restriction, we cannot achieve any reasonable approximation of IR, for the latter there exists an efficient algorithm that returns committees which approximately satisfy IR. This contrast can also be observed in our experimental results, where we analyze the attainability of IR for realistic preference profiles.

### Chapter 6: Effectiveness of Polls for Predicting Ranked Election Outcomes

Polls can build civic engagement, but public trust in polls can decline sharply when they consistently fail to predict elections. While polling for plurality elections — each voter votes for their favorite candidate

and the candidate with the most votes wins — is based on well-established statistical principles, polling for ranked-choice voting — each voter ranks the candidates and these rankings are aggregated into a collective ranking — is underexplored, despite ranked-choice voting growing in popularity and use worldwide. The question that we aim to answer in this chapter is:

> *Is it possible to make a good collective decision for a group given the ranked preferences of only a few of its members?*

In this work, we assume that there are $n$ individuals (a.k.a. voters), and we are given a voting rule $f_1$ which can make the desired collective decision given the preferences of all $n$ voters. However, we observe the preferences of only $k$ out of $n$ voters.

We propose a framework in which we are given the ranked preferences of $k$ out of $n$ individuals sampled from a distribution, and the goal is to predict what a given voting rule would output if applied on the underlying preferences of all $n$ individuals. In other words, if voting rule $f_1$ was applied on all $n$ preferences and returned a ranking, our goal is to design a voting rule $f_2$ which predicts this ranking when it is applied on the only $k$ observed preferences. We focus on the famous family of positional scoring rules, for both $f_1$ and $f_2$, which includes popular voting rules such as plurality, Borda count, harmonic rule, $k$-approval, and veto. We derive a strong negative result which indicates that even when observing all but one vote at random, if $f_1$ is any positional scoring rule, expect for plurality and veto, the probability of any positional scoring rule $f_2$ to predict the correct outcome can be zero in the worst case. We highlight that this holds for any distribution from which the $k$ observed votes are sampled. When $f_1$ is plurality (resp. veto) and we sample $k$ votes uniformly at random, plurality (resp. veto) is the best predictor, but its accuracy is still small, especially when the number of alternatives is large. These theoretical results cast a pessimistic view of the effectiveness of polls when used to predict ranked-choice elections.

We complement our theoretical results with experiments which indicate that when the voting profiles are drawn from the uniform prior, most voting rules cannot be predicted with good accuracy. However, when the profiles are drawn from a concentrated prior, the prediction accuracy of most voting rules is high, in the average case.

## Part III: Algorithmic Fairness

In this part of the thesis, we exploit notions of fairness that have been studied in economics literature by effectively applying them to different domains, like machine learning and other real life applications. In particular, we consider three traditional machine learning problems: clustering, reinforcement learning and recommendations systems, and two more real life problems: conference review assignment and team formation. We utilize famous notions of fairness from economics like envy-freeness, core and Nash social welfare. The first two are used across different domains.

A notion of group fairness that we use in the contexts of clustering, conference review assignment and team formation is inspired by the core. The core is a concept derived from cooperative game theory, where it represents the set of outcomes that cannot be improved upon by any subgroup of players. When the core is used as a notion of group fairness, the main idea is to ensure that a group of individuals collectively receives a fair share of the resources. Specifically, in the context of these problems, the core ensures that no subset of individuals should be able to find a better allocation of their entitled share of the collective resources. In contrast with other notions of fairness that have been studied mainly in the

machine learning context, such as equalized odds [Hardt et al., 2016, Berk et al., 2018] and demographic parity [Feldman et al., 2015, Zafar et al., 2017], core has the nice property that the groups are not predefined, but it provides fairness guarantee to every possible group of agents.

A notion of individual fairness that we use in the contexts of recommendation systems and team formation is based on envy-freeness. The main idea of envy-freeness is to ensure that no individual envies the allocation of another individual. In other words, if a resource is distributed among a set of individuals, envy-freeness guarantees that each person perceives their own allocation as at least as favorable as that of any other individual. In the contexts that we consider in this thesis, an allocation of an agent can be for example the team that an agent has been assigned to or the agents that the agent is matched with. In contrast, with other notions of individual fairness that have been considered in machine learning literature, like the one defined by Dwork et al. [2012], envy-freeness takes into account the preferences of the agents involved.

**Chapter 7: Proportionally Fair Clustering**

In this chapter, we focus on the centroid clustering problem, where $k$ cluster centers must be placed given $n$ points in a metric space, and the cost to each point is its distance to the nearest cluster center. We aim to find fair clustering solutions. In the last few years, many works have focused on defining reasonable notions of fairness for this problem and providing algorithms that (approximately) satisfy them [Chhabra et al., 2021].

Chen et al. [2019] introduce the notion of a *proportionally fair* clustering. The intuition of this notion is that every group of at least $n/k$ points "deserves" or is entitled to choose one cluster center. Formally, a clustering solution satisfies proportional fairness, if no group of at least $n/k$ points can find a new cluster center which provides lower cost to each member of the group. This notion of fairness is inspired by the well-known notion of group fairness in economics, called core, but to maintain consistency with previous work, we will refer to a clustering solution as proportionally fair if it satisfies this property.

Note that in contrast to most fairness definitions in machine learning which assume the existence of pre-defined groups [Backurs et al., 2019, Bera et al., 2019, Chierichetti et al., 2017, Rösner and Schmidt, 2018], this notion guarantees fairness to arbitrarily defined groups of all possible sizes.

Chen et al. [2019] propose a *greedy capture* algorithm which provides a $1 + \sqrt{2}$ approximation of proportional fairness for any metric space, and derive generalization bounds for learning proportionally fair clustering from samples in the case where a cluster center can only be placed at one of finitely many given locations in the metric space.

We focus on the case where cluster centers can be placed anywhere in the (usually infinite) metric space. In case of the $L^2$ distance metric over $\mathbb{R}^t$, we show that the approximation ratio of greedy capture improves to 2. We also show that this is due to a special property of the $L^2$ distance; for the $L^1$ and $L^\infty$ distances, the approximation ratio remains $1 + \sqrt{2}$. We provide universal lower bounds which apply to all algorithms.

We also consider metric spaces defined on graphs. For trees, we show that an exact proportionally fair clustering always exists and provide an efficient algorithm to find one. The corresponding question for general graph remains an interesting open question.

Finally, we show that for the $L^2$ distance, checking whether a proportionally fair clustering exists and implementing greedy capture over an infinite metric space are NP-hard problems, but (approximately) solvable in special cases. We also derive generalization bounds which show that an approximately

proportionally fair clustering for a large number of points can be learned from a small number of samples.

### Chapter 8: Fair Algorithms for Multi-Agent Multi-Armed Bandits

In this chapter, we propose a multi-agent variant of the classical multi-armed bandit problem. In the classical multi-armed bandit (MAB) problem, there are $K$ arms and the pulling of an arm generates a stochastic reward drawn from an unknown distribution with an unknown mean. The goal is to learn the *best arm*, which is the arm with the highest mean, by pulling sub-optimal arms as few times as possible.

Here, we consider a variation of this classic problem where the pulling of an arm generates $N$ different stochastic rewards drawn from $N$ different unknown distributions with $N$ different unknown means. This problem captures cases where making the decision of pulling an arm affects heterogeneously many agents. A concrete example is the design of a non-personalized recommendation system. Unlike the classical multi-armed bandit problem, the goal is not to learn the "best arm"; indeed, each agent may perceive a different arm to be the best for her personally. Instead, we seek to learn a fair distribution over the arms. Drawing on a long line of research in economics and computer science, we use the *Nash social welfare* as our notion of fairness.

We consider three classic algorithms for the multi-armed bandit problem: Explore-First, Epsilon-Greedy, and UCB [Slivkins et al., 2019]. We propose natural multi-agent variants of these three algorithms which take the Nash social welfare objective into account and select a distribution over the arms in each round.

Explore-First is an algorithm where the exploration and exploitation phases are well separated. In particular, this algorithm first explores all the arms by pulling each of them the same number of times, and then it exploits what it learned by choosing in all the remaining rounds the best policy according to the mean values calculated after the exploration phase. We show that this algorithm achieves a regret of $\widetilde{\mathcal{O}}\left(K^{1/3}T^{2/3}\right)$, where $T$ is the time horizon.

Epsilon-Greedy is a different algorithm, where again the phases of exploration and exploitation are separate, but now the exploration phase is spread over time. More specifically, in each round the algorithm flips a coin to decide if it will explore or exploit. The exploration probability decreases as the time passes. We show that this algorithm achieves the same regret bound as Explore-First, but its main advantage is that it does need to know the horizon $T$ upfront.

UCB is the algorithm where the phases of exploration and exploitation are not separate which means that exploration is adaptive. We show that the regret bound of the algorithm is $\widetilde{\mathcal{O}}\left(K^{1/2}T^{1/2}\right)$, where $\widetilde{\mathcal{O}}$ hides log factors. Traditionally, the horizon $T$ is often significantly larger than the number of arms $K$, thus the focus is on optimizing the regret with respect to this factor. We provide a lower bound which shows that UCB is optimal with respect to $T$.

### Chapter 9: Two-sided Markets Meets Fair Division

In this chapter, we focus on two-sided markets, and in particular, on online platforms, such as Meta, Amazon and Netflix, that are widely used for recommending products to customers. Since these platforms are profit-based, they usually focus on maximizing costumers experience. However, this practise can lead to recommendations that are largely unfair with respect to different products. For

example, popular products are exposed significantly more than products that are less popular or relatively new to the market.

Under this observation, we suggest recommendation solutions that are fair with respect to the customers and to products concurrently. We introduce a model that captures these applications which is based on two-sided matching. Specifically, we consider many-to-many matching where each agent on each side is matched to multiple agents from the other side and the utilities are additive. In this problem, traditionally a matching is considered desirable if it is stable which means that there are no two agents from different sides that are not matched but prefer each other than some of their current matches. A main drawback of this approach is that it boosts winner-take-all effects [Abdollahpouri et al., 2019], since when customers prefer a product $A$ slightly more than a product $B$ , under stability product $A$ should be exposed to all the customers by leaving less opportunities for product $B$ to be promoted.

Here, we borrow two famous fairness notions from the fair division literature namely, envy-freeness up to one good and maximin share guarantee. We demand fairness for each side separately, giving rise to notions such as double envy-freeness up to one match (DEF1) and double maximin share guarantee (DMMS).

Our main result is that when both sides have identical preferences, the famous round robin algorithm with a carefully designed agent ordering achieves (a slight strengthening of) DEF1. However, when the preferences are heterogeneous even in one side, a DEF1 matching is not guaranteed to exist. On the other hand, we show that even when both sides have identical preferences, a DMMS matching may not exist. In addition, we show that approximate DMMS and approximate DEF1 are incompatible. This indicates that the choice of the appropriate notion of fairness depends on the application. On the positive side, when the degree constraint is equal to two, i.e. each agent can be matched with up to two agents from the other side, we can achieve exact versions of both simultaneously.

## Chapter 10: Group Fairness in Peer Review

In this chapter, we consider the problem of assigning papers to reviewers in large conferences and we show how a notion of fairness that derives from economics can ensure group fairness in this context.

A main goal of the computer science conferences is to achieve high reviewing quality in order for authors of the submissions to receive useful feedback. In the past few years, conferences like AAAI, NeurIPS and ICML have grown tremendously. While on the one hand this has attracted submissions from a large number of communities, on the other hand it has also resulted in a poor reviewing experience for some communities, whose submissions end up being assigned to less qualified reviewers outside of their communities. An often-advocated solution is to break up such large conferences into smaller conferences to decentralize the reviewing process [Peng et al., 2022]. However, this can lead to isolation of various communities and slower emergence of truly interdisciplinary ideas. So, a reasonable question is:

> How can we assign submissions to reviewers in large conferences in a way that any group of researchers is fairly treated?

We tackle this challenge by introducing a notion of group fairness, namely the core, to the peer review setting. A reviewing assignment is in the core if every subset of researchers (a possible community) is treated in such a manner that they cannot achieve a better outcome by breaking off and organizing

a smaller conference on their own. We highlight that this definition provides fairness to *every* subcommunity and not only to pre-defined ones. This property is particularly useful for computer science where the last decade many interdisciplinary subcommunities have emerged and more continue to appear.

We study a simple peer review model, prove that it always admits a reviewing assignment in the core, and design an efficient algorithm to find one such assignment. On the negative side, we show that the core is incompatible with achieving a good worst-case approximation of social welfare, an often-sought desideratum. We complement these results by conducting experiments with real data. We observe that our algorithm, in addition to satisfying the core, generates good social welfare on average. In contrast, existing review assignment systems violate the core, treat many communities unfairly, and significantly incentivize them to disengage.

### Chapter 11: Fairness in Team Formation

In the last chapter, we show how two different fairness notions from economics can be applied to the real life problem of team formation for providing meaningful fairness axioms. In particular, we consider the problem of partitioning $n$ agents in an undirected social network into $k$ almost equal in size (differing by at most one) groups, where the utility of an agent for a group is the number of her neighbors in the group. Applications of this problem are the assignment of students in a course or employees in a firm to different groups, where connections in the social network indicate friendships.

The core and envy-freeness are two compelling axiomatic fairness guarantees in such settings. The former demands that there is no coalition of agents such that each agent in the coalition has more utility under the coalition than under her own group, while the latter demands that no agent envy another agent for the group she is in. While both axioms are reasonable, we show that there are instances where these axioms are incompatible. Intuitively, in some cases while the envy-freeness axiom tries to balance the number of agents from a large group of friends across different parts as much as possible, the core axiom tries to assign as many of the agents from this group into the same group. We consider the two different axioms separately.

Starting with the core, we consider bicriteria approximations of the form $(\alpha, \beta)$-core, where the first term captures a multiplicative approximation and the second term captures an additive approximation. We start with the important special case where $k = 2$ and show that $(2, 0)$-core is non-empty, and a 2-partition in the $(3, 0)$-core can be found in polynomial time. For larger $k$, we show what if we minimize the cut, we find a solution in the $(2k - 1, 0)$-core and this is the best that we can hope by this approach. While finding a partition with minimum cut is known to be NP-hard, we also present a polynomial-time algorithm that achieves the same approximation guarantee.

For envy-freeness, we also consider an additive approximation. In particular, we consider EF-$r$ which asks that non agent's utility can increase by more than $r$ if she swaps places with another agent. We make an interesting connection with discrepancy theory [Chen et al., 2014] and show that a EF-$O(\sqrt{\frac{n}{k} \cdot \log k})$ partition can be computed efficiently.

## 1.2 Bibliographic Notes

The research presented in this thesis is based on joint work with many co-authors. In all the works included, I was the primary contributor or one of multiple primary contributors. On all my publications,

the authors appear in an alphabetical order.

In Part I, Chapter 2 is based on joint work with Simon Finster, Michelle González Amador, Edwin Lock, Francisco Marmolejo-Cossío and Ariel D. Procaccia. Chapter 3 is based on joint work with Shreyas Sekar and Nisarg Shah.

In Part II, Chapter 4 is based on joint work with Soroush Ebadian, Gregory Kehne, Ariel D. Procaccia and Nisarg Shah. Chapter 5 is based on joint work with Markus Brill, Jonas Israel and Jannik Peters. Chapter 6 is based on joint work with Nisarg Shah.

In Part III, Chapter 7 is based on joint work with Nisarg Shah. Chapter 8 is based on joint work with Safwan Hossain and Nisarg Shah. Chapter 9 is based on joint work with Rupert Freeman and Nisarg Shah. Chapter 10 is based on joint work with Haris Aziz and Nisarg Shah. Chapter 11 is based on joint work with Lily Li, Aleksandar Nikolov and Nisarg Shah.

# Part I

# Optimizing Decision-Making for Social Good

# Chapter 2

# Welfare-Maximizing Pooled Testing

## 2.1 Introduction

Over the course of the ongoing COVID-19 pandemic, it has become abundantly clear that testing is an indispensable tool for combating the virus. At the same time, the pandemic has underscored the fact that this very resource can be prohibitively constrained in multiple ways (e.g. lack of reagents, trained personnel or lab equipment), making it essential to study how to systematically allocate tests for the benefit of the population at hand [Kavanagh et al., 2020, Binnicker, 2020, Dryden-Peterson et al., 2021, De Georgeo et al., 2021, Dhabaan et al., 2020, Abera et al., 2020]. In this regard, pooled testing has emerged as a promising primitive for expanding the reach of limited testing resources. In a pooled test, the samples of multiple individuals are pooled together, and a single test is applied to the mixed sample. If the test is positive, at least one individual in the pool is infected, and if negative, all individuals in the pool are healthy. Indeed, the latter case is precisely where savings in applying tests occur, as a single test can verify that multiple individuals are healthy.[1]

Ultimately, testing serves two major purposes: it prevents infections from occurring by identifying infected individuals to be isolated, and it provides a means for individuals in a population to return to in-person activities after being cleared with a negative test result. Our point of departure from prior work is the observation that different individuals have different utilities for resuming in-person activities; in a scientific setting, for example, the benefit an experimentalist derives from being able to work in their lab is typically higher than that of a theoretician who is permitted to go to the office. The goal, therefore, should be to

> *... maximize the expected welfare (overall utility) of individuals who are able to resume in-person activities.*

In more detail, we explore the scenario where a population is currently in lockdown and wishes to begin in-person activities, facilitated by a fixed budget of tests. We assume that the population is heterogeneous: each individual has their own probability of being infected and (as mentioned earlier) their own utility for being able to resume in-person activities. Tests are allocated to subsets of the population as pooled tests, and an individual is allowed to return to in-person activities (and hence earn

[1] Pooled tests can also address privacy concerns, as pooling provides a certain degree of anonymity to test takers who may not agree to be tested individually.

their corresponding utility) if any of the tests to which they are assigned results negative (proving that they are healthy). A testing regime thus obtains a certain expected welfare (overall utility) with respect to the randomized realization of infections in the population.

Our problem setting is fundamentally motivated and informed by a collaboration with the Potosinian Institute of Scientific and Technological Research (IPICYT), a higher education research institution in Mexico, with the goal of providing safe, resource-optimal alternatives to fully easing a virtual work environment. The National Laboratory of Agricultural, Medical and Environmental Biotechnology (LANBAMA), housed within IPICYT, maintains a qPCR testing facility. Three testing constraints exhibited by the LANBAMA are of crucial importance to our work: i) there are insufficient resources available to regularly test every member of the population individually, ii) the laboratory strongly prefers pooled testing protocols that do not include individuals in more than one test at a time, and iii) the laboratory conducts pooled testing on saliva samples (as opposed to nasopharyngeal sample) with a maximal pool size of $5$.

To further elucidate the latter constraint, it is important to note that although pooled testing increases resource efficiency with regards to testing reagents, this can come at a significant logistical cost for laboratory personnel when pools are overlapping, especially if the series of tests to be performed is complicated and requires delicate tab-keeping of results. In this regard, pooled testing regimes in which no individual forms part of more than one pooled test are not only conceptually simpler to study but, more importantly, logistically simpler to implement. We call such testing regimes "non-overlapping" due to the fact that pooled tests in the allocation do not overlap. By contrast, the more general overlapping testing regimes allows for the budget of tests to be allocated arbitrarily, subjecting individuals to arbitrarily many tests.

**Our results.** Despite these practical considerations that favor non-overlapping testing regimes, our first research question is whether they may be vastly outperformed by overlapping regimes. After all, if that is the case then supporting overlaps may be worth the logistical overhead. However, we show that the worst-case ratio between the welfare of the best overlapping testing regime and the best non-overlapping regime is at most 4, and in special cases even smaller. While a factor of 4 is admittedly significant, the worst example we know of gives a ratio of 7/6. Qualitatively, we interpret these results as a justification to focus on non-overlapping testing regimes. This is confirmed by our empirical results, which indicate only small gains from overlaps in practice.

Turning to the challenge of computing testing regimes, even without overlaps, the problem of determining an optimal regime is NP-hard. But we design a greedy polynomial-time algorithm that (roughly speaking) provides a 5-approximation to the optimal non-overlapping testing regime in the worst case. We then compare the performance of our greedy algorithm with optimal non-overlapping testing empirically. In our computational experiments, we choose a population size, pool size constraints and testing budgets that mirror realities at IPICYT. In order to compute (approximately) optimal non-overlapping testing regimes, we model the problem as a mixed-integer linear program (MILP). Our results indicate that the greedy algorithm computes near-optimal testing regimes, and vastly outperforms the MILP approach with respect to running time.

**Related work.** Pooled testing dates back to the seminal work of Dorfman [1943] and has since become a mature field in its own right with a rich literature of protocols typically aimed at solving the following

problem: *precisely ascertain the infection status of all individuals in a population with the minimum number of tests.* As mentioned above, our work departs significantly from this objective, as we instead assume that testing resources are initially limited, and with this provide welfare-optimizing testing allocations.

Resource constraints have motivated recent work aimed at optimally utilizing limited testing capabilities to help local communities. Ely et al. [2021] study a model where a policymaker can employ tests of different types, each with differential costs and sensitivities. The policymaker has an overall budget, and testing allocations are measured with respect to the rate at which they correctly classify individuals as infected or healthy. Brault et al. [2021] focus on limited pooled tests for early screening at a non-diagnostic level with high penalties associated with false negatives. Gollier and Gossner [2020] study pooled tests as a means to estimate infection prevalence and to find healthy individuals in a population. The main differences between their work and ours is that we consider a heterogeneous population as well as upper bounds on pool sizes imposed by lab constraints.

Most similar to our work is the work of Lock et al. [2021] and Jonnerby et al. [2020a,b], in which the authors consider a limited testing budget to be used over a heterogeneous population as a means of surveilling and containing viral spread (unlike our work which focuses on testing to find those who are healthy). The testing allocation problem is treated as a multi-objective optimization problem aimed at balancing viral spread with the overall cost of self-isolation. Although not cast as a pooled testing paper, the results of Goldberg and Rudolf [2020] can be interpreted as computing the optimal allocation of a single (arbitrarily large) pooled test to a heterogeneous population as in our model setting. The authors show that computing an optimal single test allocation is NP-hard, and they provide a fully polynomial-time approximation scheme (FPTAS) for finding an approximately optimal single test allocation; we use their FPTAS as a subroutine in our greedy algorithm. Finally, Larremore et al. [2021] and Augenblick et al. [2020] study testing frequency as a crucial factor to limiting viral spread in a pandemic environment, and how pooled testing can increase the reach of a rapid frequency testing regime when test are limited.

## 2.2 Model

Let $[n] = \{1, \ldots, n\}$ denote the population and $B \in \mathbb{N}$ be the testing budget, i.e. the number of available tests. Each individual in the population has an independent probability of infection given by $p_i \in [0, 1]$ and a utility $u_i \geqslant 0$ capturing their gain of returning to in-person activities.We also let $q_i = 1 - p_i$ denote the probability that an individual is healthy. A population instance $J$ is given by $(q_1, \ldots q_n, u_1, \ldots, u_n)$.

A single test consists of samples of a subset of the individuals, which we identify with a set $t \subseteq [n]$ of the individuals whose samples are included in the test. Test sizes are bounded by a pool size constraint $G \leqslant n$, so $|t| \leqslant G$ for all tests $t$. For convenience, we introduce the notation $q_S = \prod_{i \in S} q_i$, for any $S \subseteq [n]$, to express the probability that all individuals in $S$ are healthy; hence, $q_t$ is the probability that test $t$ is negative. A *testing regime* $T = (t_1, \ldots, t_B)$ is a collection of $B$ tests satisfying $|t_j| \leqslant G$ for each $j \in [B]$.

Individuals are allowed to return to in-person activities only if they are included in a negative test. For a given testing regime $T$, let $P_i^T$ denote the probability that $i \in [n]$ is included in some negative test $t_j \in T$. A testing regime only earns utility from individuals who return to in-person activities as a result of being in a negative test. We let $u(T)$ denote the aggregate expected utility, or *welfare*,

earned under testing regime $T$.[2] Linearity of expectation allows us to express the welfare of $T$ as $u(T) = \sum_{i\in[n]} u_i \cdot P_i^T$. In addition, we let $u(t) := u((t)) = q_t \left(\sum_{i\in t} u_i\right)$ for a single pooled test $t$. A testing regime $T$ is *optimal* (for a given population) if it maximizes welfare. Without loss of generality, we assume that $B < n$. If this is not the case, testing every person in the population individually is optimal.

**Independence of Infections.** In general it may be the case that infections in a population are correlated. However, we emphasize that our testing model is intended for a regime wherein all individuals in the given population are assumed to be in full lockdown, hence social interactions at the workplace do not contribute to potential correlation of infection for two key reasons: Either individuals who would potentially interact are forcibly at home, and hence no longer interact, or if the individuals are interacting at the workplace, it is because they are both in a negative test and hence cannot infect each other. It may be the case that colleagues interact outside of office hours, but this is not a phenomenon observed often during lockdown with our partners in Mexico.

**Non-overlapping testing regimes.** As discussed in Section 2.1, we are particularly interested in *non-overlapping* testing regimes that include each individual in at most one test. Formally, testing regime $T$ is *non-overlapping* if $t \cap t' = \emptyset$ for all distinct tests $t, t' \in T$. In general, $P_{i,T}$ can be a complicated expression due to correlation between overlapping tests. In non-overlapping testing regimes $T$, by contrast, test results are independent of one another and the welfare of $T$ is $u(T) = \sum_{t\in T} u(t)$.

**Gain of overlaps.** We are interested in quantifying the relative benefit provided by overlapping testing regimes over non-overlapping regimes, because the latter are not only conceptually simpler but also more feasible to implement in practice, as discussed in Section 2.1. Given a population instance $J$ and budget $B$, we define the *gain of overlaps* $\mathrm{gain}(B, J)$ as the ratio of the welfare of the optimal testing regime over the welfare of the optimal non-overlapping testing regime. Formally, we let $\mathcal{T}^B$ and $\widetilde{\mathcal{T}}^B$ respectively denote the space of all testing regimes and all non-overlapping testing regimes with testing budget $B$, and write

$$\mathrm{gain}(B, J) = \frac{\max_{T^*\in\mathcal{T}^B} u(T^*)}{\max_{T\in\widetilde{\mathcal{T}}^B} u(T)}.$$

The *gain of overlaps given a budget* $B$, denoted $\mathrm{gain}(B)$, is the worst-case gain over all possible instances $J$, that is, $\mathrm{gain}(B) = \sup_J \mathrm{gain}(B, J)$.

## 2.3 Theoretical Results

In this section, our goal is to provide upper bounds for the gain of overlaps. In order to develop intuition for cases in which the gain is greater than 1, we first study, as a warm-up, the case where there are only two available tests ($B = 2$), and show that the gain of overlaps is quite small. More generally, we show

[2] In the following, we will drop the term 'expected' for brevity, and assume that all welfares and utilities are determined in expectation.

that for any value of $B$, the gain is at most 4. Motivated by this result and the aforementioned practical constraints, we then focus on non-overlapping testing regimes and present a greedy algorithm that achieves a constant-factor approximation with respect to the optimal non-overlapping testing regime.

## Warm-Up: Gain of Overlaps when $B = 2$

We begin by studying the case in which $B = 2$. This case is particularly interesting for two reasons: first, we find the exact value of the gain by providing a lower bound and then showing that it is tight; second, this lower bound is the worst (largest) we know of, for any value of $B$.

**Proposition 1.** *For $B = 2$, $\mathrm{gain}(B) \geqslant 7/6$.*

*Proof.* Consider a population of three individuals $\{1, 2, 3\}$ given by $q_1 = q_2 = 1/2$, $q_3 = 1$ and $u_1 = u_2 = u_3 = 1$. We see that individuals 1 and 2 are identical, and therefore there are only four non-overlapping testing regimes up to symmetries.

- $T^1 = (\{1\}, \{2\})$ yields welfare $u(T^1) = 1$.
- $T^2 = (\{1\}, \{3\})$ yields welfare $u(T^2) = 3/2$.
- $T^3 = (\{1, 2\}, \{3\})$ yields welfare $u(T^3) = 3/2$.
- $T^4 = (\{1, 3\}, \{2\})$ yields welfare $u(T^4) = 3/2$.

Now consider the overlapping testing regime $T^* = (\{1, 3\}, \{2, 3\})$ in which individual 3 is tested twice. Then we have welfare $u(T^*) = 7/4$. □

It is of particular interest in the proof of Theorem 17 that the optimal welfare is achieved by testing the individual who is definitely healthy twice. This is intuitive, as we can test this individual many times without reducing the chances that any test is negative. In fact, in the proof of the upper bound, we use the property that the gain of overlaps is maximized when the probability that the individuals who are tested twice are healthy is maximized.

**Proposition 2.** *For $B = 2$, $\mathrm{gain}(B) \leqslant 7/6$.*

*Proof.* Suppose that given an instance, the optimal testing regime is $T^*$, where $t_1^* \cap t_2^* \neq \emptyset$. Let $A = t_1^* \setminus t_2^*$, $B = t_2^* \setminus t_1^*$ and $C = t_1^* \cap t_2^*$. Note that $t_1^* \cup t_2^* = A \cup B \cup C$. Without loss of generality, we assume that $q_A \geqslant q_B$ and with a slight abuse of notation, we denote $u_A = \sum_{i \in A} u_i$, $u_B = \sum_{i \in B} u_i$, and $u_C = \sum_{i \in C} u_i$. Moreover, we define the following four different testing regimes:

- $T^1$ with $t_1^1 = A \cup C$ and $t_2^1 = B$
- $T^2$ with $t_1^2 = A$ and $t_2^2 = C$
- $T^3$ with $t_1^3 = B$ and $t_2^3 = C$
- $T^4$ with $t_1^4 = A \cup B$ and $t_2^4 = C$

We start with the following necessary lemma.

**Lemma 1.** *For any $\widehat{T} \in \{T^1, T^2, T^3, T^4\}$, the ratio $u(T^*)/u(\widehat{T})$ is maximized when $q_C = 1$.*

*Proof.* We first show that the statement is true for $\widehat{T} = T^1$. We need to show that

$$\begin{aligned}
&\frac{q_C\left(q_A \cdot u_A + q_B \cdot u_B + (q_A + (1-q_A)\cdot q_B)\cdot u_C\right)}{q_C \cdot q_A \cdot u_A + q_B \cdot u_B + q_C \cdot q_A \cdot u_C} \\
&\leqslant \frac{q_A \cdot u_A + q_B \cdot u_B + (q_A + (1-q_A)\cdot q_B)\cdot u_C}{q_A \cdot u_A + q_B \cdot u_B + q_A \cdot u_C} \\
\Rightarrow q_C &\leqslant \frac{q_C \cdot q_A \cdot u_A + q_B \cdot u_B + q_C \cdot q_A \cdot u_C}{q_A \cdot u_A + q_B \cdot u_B + q_A \cdot u_C}
\end{aligned}$$

which is true since

$$q_C = q_C \cdot \frac{q_A \cdot u_A + q_B \cdot u_B + q_A \cdot u_C}{q_A \cdot u_A + q_B \cdot u_B + q_A \cdot u_C} \leqslant \frac{q_C \cdot q_A \cdot u_A + q_B \cdot u_B + q_C \cdot q_A \cdot u_C}{q_A \cdot u_A + q_B \cdot u_B + q_A \cdot u_C}.$$

Now, we consider the case that $\widehat{T} = T^2$. Here, we need to show that

$$\begin{aligned}
&\frac{q_C\left(q_A \cdot u_A + q_B \cdot u_B + (q_A + (1-q_A)\cdot q_B)\cdot u_C\right)}{q_A \cdot u_A + q_C \cdot u_C} \\
&\leqslant \frac{q_A \cdot u_A + q_B \cdot u_B + (q_A + (1-q_A)\cdot q_B)\cdot u_C}{q_A \cdot u_A + u_C} \\
\Rightarrow q_C &\leqslant \frac{q_A \cdot u_A + q_C \cdot u_C}{q_A \cdot u_A + u_C},
\end{aligned}$$

which is true since

$$q_C = q_C \cdot \frac{q_A \cdot u_A + u_C}{q_A \cdot u_A + u_C} \leqslant \frac{q_A \cdot u_A + q_C \cdot u_C}{q_A \cdot u_A + u_C}.$$

With similar arguments as above, we can show that the ratio $u(T^*)/u(T^3)$ is maximized when $q_C = 1$.

Lastly for $\widehat{T} = T^4$, we need to show that

$$\begin{aligned}
&\frac{q_C\left(q_A \cdot u_A + q_B \cdot u_B + (q_A + (1-q_A)\cdot q_B)\cdot u_C\right)}{q_A \cdot q_B \cdot (u_A + u_B) + q_C \cdot u_C} \\
&\leqslant \frac{q_A \cdot u_A + q_B \cdot u_B + (q_A + (1-q_A)\cdot q_B)\cdot u_C}{q_A \cdot q_B \cdot (u_A + u_B) + u_C} \\
\Rightarrow q_C &\leqslant \frac{q_A \cdot q_B \cdot (u_A + u_B) + q_C \cdot u_C}{q_A \cdot q_B \cdot (u_A + u_B) + u_C},
\end{aligned}$$

which is true since

$$q_C = q_C \cdot \frac{q_A \cdot q_B \cdot (u_A + u_B) + u_C}{q_A \cdot q_B \cdot (u_A + u_B) + u_C} \leqslant \frac{q_A \cdot q_B \cdot (u_A + u_B) + q_C \cdot u_C}{q_A \cdot q_B \cdot (u_A + u_B) + u_C}.$$

□

Using Lemma 1, hereinafter, we consider the case that $q_C = 1$.

We distinguish into two cases.

**Case I:** $q_A \geqslant 5/6$**.** In this case, note that

$$\frac{(q_A + (1-q_A)\cdot q_B)\cdot u_C}{q_A \cdot u_C} \leqslant \frac{(q_A + (1-q_A)\cdot q_A)\cdot u_C}{q_A \cdot u_C} = 2 - q_A \leqslant \frac{7}{6}.$$

where the second transition follows since $q_A \geqslant q_B$ and the last transition follows since $q_A \geqslant 5/6$. Hence, we see that

$$\frac{q_A \cdot u_A + q_B \cdot u_B + (q_A + (1-q_A)\cdot q_A)\cdot u_C}{q_A \cdot u_A + q_B \cdot u_B + q_A \cdot u_C} \leqslant \frac{7}{6},$$

meaning that $u(T^*)/u(T^1) \leqslant 7/6$.

**Case II:** $q_A < 5/6$**.** Here, for the sake of contradiction, suppose that for any testing without overlaps the approximation ratio is more that $7/6$. Then, we have that $u(T^*)/u(\widehat{T}) > 7/6$, for any $\widehat{T} \in \{T^1, T^2, T^3, T^4\}$.

Thus,

$$\begin{aligned}
&\frac{u(T^*)}{u(T^1)} > \frac{7}{6} \\
\Rightarrow &\frac{q_A \cdot u_A + q_B \cdot u_B + (q_A + (1-q_A)\cdot q_B)\cdot u_C}{q_A \cdot u_A + q_B \cdot u_B + q_A \cdot u_C} > \frac{7}{6} \\
\Rightarrow &(6\cdot(q_A + (1-q_A)\cdot q_B) - 7q_A)\cdot u_C > q_A \cdot u_A + q_B \cdot u_B. \qquad (2.1)
\end{aligned}$$

Now, from the fact that $u(T^*)/u(T^2) > 7/6$, we get that

$$\begin{aligned}
\Rightarrow &\frac{q_A \cdot u_A + q_B \cdot u_B + (q_A + (1-q_A)\cdot q_B)\cdot u_C}{q_A \cdot u_A + u_C} > \frac{7}{6} \\
\Rightarrow &q_B \cdot u_B > \frac{1}{6} q_A \cdot u_A + (\frac{7}{6} - (q_A + (1-q_A)\cdot q_B))\cdot u_C \\
\Rightarrow &q_B \cdot u_B > \frac{1}{6} q_A \cdot u_A + (\frac{7}{6} - (q_A + (1-q_A)\cdot q_A))\cdot u_C \qquad (2.2)
\end{aligned}$$

where the last inequality follows from the fact that $q_A \geqslant q_B$. With a very similar argument we can conclude that

$$q_A \cdot u_A > \frac{1}{6} q_B \cdot u_B + (\frac{7}{6} - (q_A + (1-q_A)\cdot q_A))\cdot u_C \qquad (2.3)$$

from the fact that $u(T^*)/u(T^3) > 7/6$. From Equation (2.2) and Equation (2.3), we get that

$$q_A \cdot u_A + q_B \cdot u_B > \frac{6}{5}\cdot 2 \cdot (\frac{7}{6} - (q_A + (1-q_A)\cdot q_A))\cdot u_C$$

and from Equation (2.1), we conclude that

$$(6\cdot(q_A + (1-q_A)\cdot q_B) - 7q_A) > \frac{6}{5}\cdot 2 \cdot (\frac{7}{6} - (q_A + (1-q_A)\cdot q_A))$$

which is true when $1/2 < q_A < 2/3$. Hence, from now one we assume that $q_A$ lies in the interval $(1/2, 2/3)$.

Moreover, we have that

$$\begin{aligned}
&\frac{u(T^*)}{u(T^2)} > \frac{7}{6} \\
\Rightarrow &\frac{q_A \cdot u_A + q_B \cdot u_B + (q_A + (1-q_A) \cdot q_B) \cdot u_C}{q_A \cdot u_A + u_C} > \frac{7}{6} \\
\Rightarrow &\frac{(6 \cdot (q_A + (1-q_A) \cdot q_B) - 7q_A) \cdot u_C + (q_A + (1-q_A) \cdot q_B) \cdot u_C}{q_A \cdot u_A + u_C} > \frac{7}{6} \\
\Rightarrow &\frac{7 \cdot (1-q_A) \cdot q_B \cdot u_C}{q_A \cdot u_A + u_C} > \frac{7}{6} \\
\Rightarrow &(6 \cdot (1-q_A) \cdot q_B - 1) \cdot u_C > q_A \cdot u_A
\end{aligned} \tag{2.4}$$

where the third inequality follows from Equation (2.1), and similarly using the fact that $u(T^*)/u(T^3)$, we get

$$(6 \cdot (1-q_A) \cdot q_B - 1) \cdot u_C > q_B \cdot u_B \tag{2.5}$$

Lastly,

$$\begin{aligned}
&\frac{u(T^*)}{u(T^4)} > \frac{7}{6} \\
\Rightarrow &\frac{q_A \cdot u_A + q_B \cdot u_B + (q_A + (1-q_A) \cdot q_B) \cdot u_C}{q_A \cdot q_B \cdot (u_A + u_B) + u_C} > \frac{7}{6} \\
\Rightarrow &q_A \cdot u_A + q_B \cdot u_B > \frac{(7 - 6(q_A + (1-q_A) \cdot q_B)) \cdot u_C}{6 - 7 \cdot q_B}
\end{aligned} \tag{2.6}$$

where the last inequality follows from the fact that $q_B < 6/7$ since $q_B \leqslant q_A$ and $q_A < 5/6$.

Now, from Equation (2.4), Equation (2.5) and Equation (2.6), we get that

$$\begin{aligned}
\frac{u(T^*)}{u(T^1)} &= \frac{q_A \cdot u_A + q_B \cdot u_B + (q_A + (1-q_A) \cdot q_B) \cdot u_C}{q_A \cdot u_A + q_B \cdot u_B + q_A \cdot u_C} \\
&\leqslant \frac{2 \cdot (6 \cdot (1-q_A) \cdot q_B - 1) \cdot u_C + (q_A + (1-q_A) \cdot q_B) \cdot u_C}{\frac{7 - 6(q_A + (1-q_A) \cdot q_B) \cdot u_C}{6 - 7 \cdot q_B} + q_A \cdot u_C}
\end{aligned}$$

which is maximized when $q_A = q_B = 1/2$, given that $1/2 < q_A < 2/3$ and then we get that $u(T^*)/u(T^1) \leqslant 7/6$ and reach a contradiction. □

## Upper Bound on Gain of Overlaps for Any $B \geqslant 2$

In this section, we show that the gain of overlaps is a small constant; not only for the case that there are two tests available, but for any $B$. Before doing so, we start with some necessary notation. Given a testing regime $T = (t_1, \ldots, t_B)$ and individual $i \in [n]$, we let $T(i) = \{t_j \in T \mid i \in t_j\}$.

Furthermore, we denote with $T(i;j)$ the set of tests with index less than $j$ in which $i$ has been tested, i.e. $T(i;j) = \{t_{j'} \in T(i) : j' < j\}$. We say that test $t_j$ is *pivotal* for individual $i$ if: $i$ is included in $t_j$, the result of $t_j$ is negative, and all tests in $T(i;j)$ are positive. Equivalently, test $t_j$ is pivotal for individual $i$ if it is the negative test of smallest index in $T(i)$. We let $P^T_{i,j}$ denote the probability that $t_j$ is pivotal for individual $i$ under random infection realizations. In other words:

$$P^T_{i,j} = \begin{cases} \Pr[\,\forall t_{j'} \in T(i;j), t_{j'} \text{ is positive and } t_j \text{ is negative}] & \text{if } t_j \in T(i), \\ 0 & \text{otherwise.} \end{cases}$$

Individual $i$ is in a negative test if and only if a single test $t_j \in T$ is pivotal for $i$, hence $P^T_i = \sum_{j\in[B]} P^T_{i,j}$. As previously advertized, our main result of this section is that overlapping testing regimes have bounded gain over non-overlapping regimes.

**Theorem 1.** *For any $B \geqslant 1$, $\mathrm{gain}(B) \leqslant 4$.*

To prove the theorem, we will show that given an optimal overlapping testing regime $T^*$, we can find a non-overlapping testing regime $T$ such that $u(T^*)/u(T) \leqslant 4$ in polynomial time. This does not lead to a polynomial-time algorithm, as it requires access to $T^*$ to begin with.

The proof requires a two lemmas. The first lemma is more intuitive: There exists an optimal testing regime $T^* = (t^*_1, \ldots, t^*_B)$ such that if $t^*_j \in T^*(i)$, it must be the case that the probability that $t^*_j$ is pivotal for $i$ is positive.

**Lemma 2.** *There exists an optimal $T^* = (t^*_1, \ldots, t^*_B)$ such that if $t^*_j \in T^*(i)$, then $P^{T^*}_{i,j} > 0$.*

*Proof.* We re-write $P^{T^*}_{i,j}$ for $i \in t^*_j$ using conditional probabilities:

$$\begin{aligned} P^{T^*}_{i,j} &= \Pr[\,\forall t^*_{j'} \in T^*(i;j), t^*_{j'} \text{ is positive and } t^*_j \text{ is negative}]. \\ &= \Pr[t^*_j \text{ negative}] \cdot \Pr[\forall t^*_{j'} \in T^*(i;j), t^*_{j'} \text{ positive} \mid t^*_j \text{ negative}] \\ &= q_{t^*_j} \cdot \Pr[\forall t^*_{j'} \in T^*(i;j)\ \exists i' \in t^*_{j'} \setminus t^*_j \mid i \text{ infected}]. \end{aligned}$$

First of all, we notice that it must be the case that $q_{t^*_j} = \prod_{i\in t^*_j} q_i > 0$. If this is not so then there must be some individual $i \in t^*_j$ such that $p_i = 1$, and it is straightforward to see that it is always sub-optimal to include such an individual in any testing regime. As for the second term, to show that it is non-zero, we begin by using the fact that $T^*$ is optimal to show that without loss of generality, we can assume that for each $t^*_{j'} \in T^*(i;j)$ it must be the case that there exists an $i \in t^*_{j'} \setminus t^*_j$ such that $p_i > 0$.

Suppose that this is not the case and that there exist $t^*_{j'}, t^*_j \in T^*$ such that every $i \in t^*_{j'} \setminus t^*_j$ has $p_i = 0$ ($q_i = 1$). We show that either $T^*$ is sub-optimal, or we can construct an optimal testing regime where this is no longer the case. To do so, we assume that without loss of generality $j' = 1$ and $j = 2$ (we can arbitrarily re-order test indices), and write the expected utility of $T^*$ as follows:

$$
\begin{aligned}
u(T^*) &= \sum_{i\in[n]} u_i \cdot P_i^{T^*} \\
&= \sum_{i\in[n]} \sum_{j\in[B]} u_i \cdot P_{i,j}^{T^*} \\
&= \sum_{j\in[B]} \sum_{i\in t_j^*} u_i \cdot P_{i,j}^{T^*} \\
&= \left(\sum_{i\in t_1^*} u_i \cdot P_{i,1}^{T^*}\right) + \left(\sum_{i\in t_2^*} u_i \cdot P_{i,2}^{T^*}\right) + \sum_{j=3}^{B} \sum_{i\in t_j^*} u_i \cdot P_{i,j}^{T^*} \\
&= u(t_1^*) + \left(\sum_{i\in t_1^*} u_i \cdot P_{i,2}^{T^*} + \sum_{i\in t_2^*\setminus t_1^*} u_i \cdot P_{i,2}^{T^*}\right) + \sum_{j=3}^{B} \sum_{i\in t_j^*} u_i \cdot P_{i,j}^{T^*} \\
&= u(t_1^*) + \left(\sum_{i\in t_2^*\setminus t_1^*} u_i \cdot P_{i,2}^{T^*}\right) + \sum_{j=3}^{B} \sum_{i\in t_j^*} u_i \cdot P_{i,j}^{T^*} \\
&= u(t_1^*) + q_{t_2^*} \sum_{i\in t_2^*\setminus t_1^*} u_i + \sum_{j=3}^{B} \sum_{i\in t_j^*} u_i \cdot P_{i,j}^{T^*} \\
&\leqslant u(t_1^*) + u(t_2^* \setminus t_1^*) + \sum_{j=3}^{B} \sum_{i\in t_j^*} u_i \cdot P_{i,j}^{T^*}
\end{aligned}
$$

The first 4 equalities follow from re-ordering terms in the sums. In the fifth equality, we make use of the fact that if $i \in t_1^*$, it must be the case that $P_{i,2}^{T^*} = 0$, for if $t_2^*$ is pivotal for them, then that test must be negative, which implies that $t_1^*$ is negative (for individuals in $t^* \setminus t_2^*$ are guaranteed to be healthy by assumption), contradicting the pivotal nature of $t_2^*$. If $i \in t_2^* \setminus t_1^*$, then $t_2^*$ is pivotal only if it is negative, hence $P_{i,2}^{T^*} = q_{t_2^*}$, which in turn justifies the following equality. Finally, we know that $q_{t_2^*} \leqslant q_{t_2^*\setminus t_1^*}$, and that $u(t_2^* \setminus t_1^*) = q_{t_2^*\setminus t_1^*} \sum_{i\in t_2^*\setminus t_1^*} u_i$, from which the final equality holds. Putting everything together, let us consider $T'$ where $t'_j = t_j^*$ for $j \neq 2$ and $t'_2 = t_2^* \setminus t_1^*$. From the above, it follows that either $T^*$ is sub-optimal (as $T'$ achieves more welfare), or we can replace $T^*$ with $T'$, and in either case, we can ensure that our desired property holds.

With this in hand, this means if $T^*$ is optimal, and we consider $i \in t_j^*$ as in the beginning of the proof, we can construct a pool of individuals $S$ by picking one individual from the set $T_j^* \setminus T_{j'}^*$ for each $T_{j'}^* \in T(i;j)$ such that the individual has non-zero probability of infection. From the above, we are guaranteed to be able to construct such an $S$ which is non-empty. Furthermore, if all individuals in $S$ are infected, it follows that each $T_{j'} \in T(i;j)$ is positive without compromising a negative test on $T_j^*$. This in turn implies that $\Pr[\forall t_{j'}^* \in T^*(i;j)\ \exists i' \in t_{j'}^* \setminus t_j^* \mid i \text{ infected}] \geqslant \Pr[1 - q_S] > 0$. Putting this together with the fact that $q_{t_j^*} > 0$ completes the proof. □

The second lemma is a non-trivial generalization of Lemma 6 of Goldberg and Rudolf [2020]. At a high level, we show that if a (non-overlapping) testing regime is optimal, no test within this regime

can be split into two groups which simultaneously have a "low" probability of being healthy. When the testing regime is non-overlapping, the generalization is straightforward, but for the general case where the testing regime may be overlapping, novel techniques and arguments are used.

**Lemma 3.** *Suppose that $T^* = (t_1^*, \ldots, t_B^*)$ is an optimal (non-overlapping) testing regime and that $\alpha \in (0,1)$. For any $t_j^*$ and any $S \subset t_j^*$, if $q_S < \alpha$, then $q_{t_j^* \setminus S} \geqslant 1 - \alpha$.*

*Proof.* First note that for any $i \in t_j^*$,

$$\begin{aligned}
P_{i,j}^{T^*} &= \Pr[\,\forall t_{j'}^* \in T^*(i;j), t_{j'}^* \text{ is positive and } t_j^* \text{ is negative}] \\
&= \Pr[\,\forall t_{j'}^* \in T^*(i;j), t_{j'}^* \text{ is positive } \mid t_j^* \text{ is negative}] \cdot \Pr[t_j^* \text{ is negative}] \\
&= \Pr[\,\forall t_{j'}^* \in T^*(i;j), t_{j'}^* \setminus t_j^* \text{ is positive } \mid t_j^* \text{ is negative}] \cdot \Pr[t_j^* \text{ is negative}] \\
&= \Pr[\,\forall t_{j'}^* \in T^*(i;j), t_{j'}^* \setminus t_j^* \text{ is positive }] \cdot \Pr[t_j^* \text{ is negative}] \\
&= \Pr[\,\forall t_{j'}^* \in T^*(i;j), t_{j'}^* \setminus t_j^* \text{ is positive }] \cdot q_{t_j^*}
\end{aligned}$$

where the third equality follows since $t_{j'}^*$ can be positive if and only if some individual in $t_{j'}^* \setminus t_j^*$ is infected as given that $t_j^*$ is negative we conclude that any individual in $t_j^*$ is healthy, and the fourth inequality follows since each individual has an independent probability to be infected.

Assume for the sake of contradiction that for some $j \in [B]$, there exists $S \subset T_{j^*}^*$ such that $q_S < \alpha$ and $q_{T_{j^*}^* \setminus S} < 1 - \alpha$. Without loss of generality, we assume that $j^* = B$. Then, since from Lemma 2 we know that $P_{i,j}^{t_j^*} > 0$ for any individual $i$ such that $i \in T_j^*$, we have that

$$\begin{aligned}
u(T^*) &= \sum_{i \in [n]} P_i^T \cdot u_i \\
&= \sum_{i \in [n]} \sum_{j \in [B]} P_{i,j}^{T^*} \cdot u_i \\
&= \sum_{j \in [B-1]} \sum_{i \in [n]} P_{i,j}^{T^*} \cdot u_i + \sum_{i \in [n]} I_{i,B}^{T^*} \cdot q_{t_B^*} \cdot \Pr[\forall t_{j'}^* \in T^*(i;B), t_{j'}^* \setminus t_B^* \text{ is positive}] \cdot u_i \\
&= \sum_{j \in [B-1]} \sum_{i \in [n]} P_{i,j}^{T^*} \cdot u_i + q_S \cdot q_{t_B^* \setminus S} \cdot \sum_{i \in t_B^*} \Pr[\forall t_{j'}^* \in T^*(i;B), t_{j'}^* \setminus t_B^* \text{ is positive}] \cdot u_i \\
&= \sum_{j \in [B-1]} \sum_{i \in [n]} P_{i,j}^{T^*} \cdot u_i + q_{t_B^* \setminus S} \cdot \left( q_S \cdot \sum_{i \in S} \Pr[\forall t_{j'}^* \in T^*(i;B), t_{j'}^* \setminus t_B^* \text{ is positive}] \cdot u_i \right) \\
&\quad + q_S \left( q_{t_B^* \setminus S} \cdot \sum_{i \in t_B^* \setminus S} \Pr[\forall t_{j'}^* \in T^*(i;B), t_{j'}^* \setminus t_B^* \text{ is positive}] \cdot u_i \right) \\
&\leqslant \sum_{j \in [B-1]} \sum_{i \in [n]} P_{i,j}^{T^*} \cdot u_i + q_{t_B^* \setminus S} \cdot \left( q_S \cdot \sum_{i \in S} \Pr[\forall t_{j'}^* \in T^*(i;B), t_{j'}^* \setminus S \text{ is positive}] \cdot u_i \right) \\
&\quad + q_S \left( q_{t_B^* \setminus S} \cdot \sum_{i \in t_B^* \setminus S} \Pr[\forall t_{j'}^* \in T^*(i;B), t_{j'}^* \setminus (t_B^* \setminus S) \text{ is positive}] \cdot u_i \right)
\end{aligned}$$

$$< \sum_{j\in[B-1]}\sum_{i\in[n]} P^{T^*}_{i,j}\cdot u_i + \max\Bigg\{ q_S \cdot \sum_{i\in S} \Pr[\forall t^*_{j'} \in T^*(i;B), t^*_{j'} \setminus S \text{ is positive}]\cdot u_i,$$

$$q_{t^*_B\setminus S}\cdot \sum_{i\in t^*_j\setminus S} \Pr[\forall t^*_{j'} \in T^*(i;B), t^*_{j'} \setminus (t^*_B \setminus S) \text{ is positive}]\cdot u_i \Bigg\}$$

To justify the first inequality, we begin by showing that

$$\sum_{i\in S} \Pr[\forall t^*_{j'} \in T^*(i;j),\ t^*_{j'} \setminus t^*_B \text{ is positive}] \leqslant \sum_{i\in S} \Pr[\forall t^*_{j'} \in T^*(i;j),\ t^*_{j'} \setminus S \text{ is positive}].$$

Since $S \subseteq t^*_B$, it follows that for any $t^*_{j'} \in T^*(i;j)$ we have $t^*_{j'} \setminus t^*_B \subseteq t^*_{j'} \setminus S$. This in turn implies that if $t^*_{j'} \setminus t^*_B$ is positive, it must be the case that $t^*_{j'} \setminus S$ is positive. It follows that the event where $\forall t^*_{j'} \in T^*(i:j),\ t^*_{j'} \setminus t^*_B$ is positive implies that $\forall t^*_{j'} \in T^*(i:j),\ t^*_{j'} \setminus S$ is positive, hence the inequality follows. Furthermore, the argumentation above can be replicated with $t^*_B \setminus S$ rather than $S$ to fully justify the inequality from the main derivation. As for the second inequality, in the derivation, this follows from the assumption that $q_S < \alpha$ and $q_{t^*_B\setminus S} < 1-\alpha$.

To reach a contradiction, let $T$ be a testing regime with $t_j = t^*_j$ for any $j \in [B-1]$ and $t_B = S$ if

$$\begin{aligned} & q_S \cdot \sum_{i\in S} \Pr[\forall t^*_{j'} \in T^*(i;B), t^*_{j'} \setminus S \text{ is positive}]\cdot u_i \\ & \geqslant q_{T^*_B\setminus S}\cdot \sum_{i\in T^*_B\setminus S} \Pr[\forall t^*_{j'} \in T^*(i;B), t^*_{j'} \setminus (t^*_B \setminus S) \text{ is positive}]\cdot u_i \end{aligned}$$

and $t_B = t^*_B \setminus S$, otherwise. Then,

$$\begin{aligned} u(T) &= \sum_{i\in[n]} P^T_i \cdot u_i \\ &= \sum_{i\in[n]}\sum_{j\in[B]} P^T_{i,j}\cdot u_i \\ &= \sum_{j\in[B-1]}\sum_{i\in[n]} P^{T^*}_{i,j}\cdot u_i + \sum_{i\in t_B} q_{t_B}\cdot \Pr[\forall t^*_{j'} \in T^*(i;B), t^*_{j'} \setminus t_B \text{ is positive}]\cdot u_i \\ &= \sum_{j\in[B-1]}\sum_{i\in N} P^{T^*}_{i,j}\cdot u_i + \max\Bigg\{ q_S \cdot \sum_{i\in S} \Pr[\forall t^*_{j'} \in T^*(i,B), t^*_{j'} \setminus S \text{ is positive}]\cdot u_i, \\ &\qquad q_{t^*_B\setminus S}\cdot \sum_{i\in T_B\setminus S} \Pr[\forall t^*_{j'} \in T^*(i;B), t^*_{j'} \setminus (t^*_B \setminus S) \text{ is positive}]\cdot u_i \Bigg\} \\ &> u(T^*), \end{aligned}$$

which contradicts the optimality of $T^*$ and completes the proof of our claim. □

We are now ready to prove the theorem.

*Proof of Theorem 1.* We begin by constructing an intermediate non-overlapping testing regime $T$ from $T^*$ by simply assigning each individual to only a single test chosen arbitrary among all tests to which that individual is assigned to in $T^*$. Thus, we have that for each $j \in [B]$ $t_j \subseteq t_j^*$ which meana that $q_{t_j} \geqslant q_{t_j^*}$. Now, for each $t_j$, we let $S_j$ be the smallest subset of $t_j$ such that $q_{S_j} < 1/2$ (if $q_{t_j} \geqslant 1/2$, then $S_j = \emptyset$). Note that for each $i \in S_j$, we have $q_{S_j \setminus \{i\}} \geqslant 1/2$, as $S_j$ is the smallest possible set that has probability less than $1/2$ to be negative. In addition, we can show that $q_{t_j \setminus S_j} \geqslant 1/2$. To see this, notice that $S_j \subseteq t_j \subseteq t_j^*$ and $q_{S_j} < 1/2$. From Lemma 3, we know that $q_{t_j^* \setminus S_j} \geqslant 1/2$, but since $t_j \setminus S_j \subseteq t_j^* \setminus S_j$, it follows that $q_{t_j \setminus S_j} \geqslant q_{t_j^* \setminus S_j} \geqslant 1/2$ as desired.

Next, consider two disjoint testing regimes given by $T^1$ where $t_j^1 = S_j$ and $T^2$ where $t_j^2 = t_j \setminus S_j$. Using the properties of $T$ from the previous paragraph, we wish to show that for each $i \in t_j^\ell$ where $\ell \in \{1, 2\}$, we have that $P_i^{T^\ell} \geqslant q_i \cdot 1/2$. To that end, in the case of $\ell = 1$, we get $P_i^{T^1} = q_{t_j^1} = q_{S_j} < 1/2$. However, we also know that $q_{t_j^1 \setminus \{i\}} = q_{S_j \setminus \{i\}} \geqslant 1/2$, and hence $q_{t_j^1} = q_i \cdot q_{t_j^1 \setminus \{i\}} \geqslant q_i \cdot 1/2$ as desired. As for the case where $\ell = 2$, we get $P_i^{T^2} = q_{t_j^2} = q_{t_j \setminus S_j}$. The right hand side can be decomposed as $q_{t_j \setminus S_j} = q_i \cdot q_{t_j^2 \setminus \{i\}}$. However, as we have shown above, the choice of $S_j$ ensures that $q_{t_j \setminus S_j} \geqslant 1/2$, and it follows that $q_i \cdot q_{t_j^2 \setminus \{i\}} \geqslant 1/2 \geqslant q_i \cdot 1/2$. We conclude that $P_i^{T^2} \geqslant q_i \cdot 1/2$, as desired.

Without loss of generality, assume that $\bigcup_{j \in [B]} t_j^* = [n']$ for some $n' \leqslant n$, i.e., the first $n'$ individuals are included in some test under $T^*$. Notice that $\bigcup_{j \in [B]} t_j = [n']$, as each individual who is included in some test under $T^*$ is also included in some test under $T$, and no individual who is not included in some test under $T^*$ is included in $T$. Thus, we get that

$$u(T^*) = \sum_{i \in [n']} P_i^{T^*} \cdot u_i = \sum_{j \in [B]} \left( \sum_{i \in S_j} P_i^{T^*} \cdot u_i + \sum_{i \in T_j \setminus S_j} P_i^{T^*} \cdot u_i \right).$$

We now consider two cases, depending on whether $\sum_{j \in [B]} \sum_{i \in S_j} P_i^{T^*} \cdot u_i \geqslant \sum_{j \in [B]} \sum_{i \in t_j \setminus S_j} P_i^{T^*} \cdot u_i$. If this is the case, we get that

$$u(T^*) \leqslant 2 \cdot \sum_{j \in [B]} \sum_{i \in S_j} P_i^{T^*} \cdot u_i \leqslant 2 \cdot \sum_{j \in [B]} \sum_{i \in S_j} q_i \cdot u_i.$$

It also holds that

$$u(T^1) = \sum_{j \in [B]} \sum_{i \in S_j} P_i^{T^1} \cdot u_i \geqslant \sum_{j \in [B]} \sum_{i \in S_j} 1/2 \cdot q_i \cdot u_i.$$

Thus, we conclude that

$$\frac{u(T^*)}{u(T^1)} \leqslant \frac{2 \sum_{i \in S_j} q_i \cdot u_i}{\sum_{i \in S_j} 1/2 \cdot q_i \cdot u_i} \leqslant 4.$$

Using similar arguments, when $\sum_{j \in [B]} \sum_{i \in S_j} P_i^{T^*} \cdot u_i < \sum_{j \in [B]} \sum_{i \in T_j \setminus S_j} P_i^{T^*} \cdot u_i$, we can show that $u(T^*)/u(T^2) \leqslant 4$ and the theorem follows. □

While this proves that the gain of overlaps cannot be larger than $4$, the worst known example is the one illustrated in Theorem 17, providing a lower bound of $7/6$. Interestingly, after running many simulations we were not able to find a better lower bound, and we believe that the gain of overlaps is less than $4$.

## Greedy Algorithm for the Non-Overlapping Testing Regime

In light of the previous result, hereinafter we focus on non-overlapping testing regimes. Consider the case where $B = 1$. If $G$ is a constant, we can efficiently enumerate all $O(n^G)$ potential pooled tests and find the optimal test $t^*$. On the other hand, when $G = n$, it follows from the work of Goldberg and Rudolf [2020] that even when there is only one test, it is NP-hard to find the subset of individuals that maximizes the expected welfare of the test. On the positive side, they provide a fully polynomial-time approximation scheme (FPTAS) for the same case. Here, we show that we can adjust the main ideas of their algorithm to obtain an FPTAS for the case where there is one test with up to $G \in [n]$ individuals in it.

**Lemma 4.** *When $B = 1$, there is an FPTAS for computing approximately optimal $(t^*) \in \mathcal{T}^1$.*

*Proof.* Here we show, how the FPTAS, that is introduced in [Goldberg and Rudolf, 2020] and finds an almost optimal test when there is no any size constraint, can be modified for the case that a test can pool up to $G$ samples.

Using similar notation as in [Goldberg and Rudolf, 2020], for $i \in [n]$, we denote with $P(i, C, L)$ the maximum probability of a subset of $[i]$ to be negative with sum of utilities exactly $C$ and size exactly $L$. Then, we modify the dynamic program that was introduced at Equation (6) in [Goldberg and Rudolf, 2020] as following:

$$P(i,C,L) = \begin{cases} \max\{P(i-1,C,L), q_i \cdot P(i-1,C-u_i,L-1)\} & i \geqslant 2 \text{ and } u_i < C \\ P(i-1,C,L) & i \geqslant 2 \text{ and } u_i \geqslant C \\ q_i & i = 1 \text{ and } u_1 = C \\ 0 & \text{otherwise} \end{cases} \tag{2.7}$$

With a slight abuse of notation, we denote with $t(P(i,C,L))$ the subset of $[i]$ that satisfies $q_{t(P(i,C,L))} = P(i,C,L)$, $\sum_{\ell \in [t(P(i,C,L))]} u_\ell = C$ and $|t(P(i,C,L))| = L$.

If $\overline{C}$ is an upper bound on the sum of utilities of the optimal test ( a straightforward upper bound is $\overline{C} = \sum_{i\in[n]} u_i$), then the optimal test with size at most $G$ that maximizes the utility is given by $t(P(n, C^*, L^*))$ where

$$C^*, L^* = \underset{C\in[\overline{C}], L\in[G]}{\arg\max}\ C \cdot P(n,C,L). \tag{2.8}$$

Thus, the running time is given by $O(nG\overline{C}) \leqslant O(n^2\overline{C})$, since $G \leqslant n$. Then, we see that in order to approximately solve the dynamic programming with a polynomial run-time complexity bound, we should scale down (and round) the utility coefficients whose magnitude determines the running time of the program. We can achieve this by using identical arguments as in Section 3.2 of [Goldberg and Rudolf, 2020]. We present the whole proof here for completeness.

We scale down the utilities using some factor $\kappa$, by setting $\widehat{u}_i = \lfloor u_i/\kappa \rfloor$ for each $i \in [n]$. Before, we choose $\kappa$, we add some more notation. Let $N_{1/2} = \{i \in [n] : q_i \geqslant 1/2\}$. Without loss of generality, assume that $N_{1/2} = [h]$ and $[n] \setminus N_{1/2} = \{h+1, \ldots, n\}$. Let $\widehat{P}(i,C,L)$ denote the DP in Equation (2.7) by replacing $u_i$ with $\widehat{u}_i$. Moreover, we assume that there exists a dummy individual $n+1$ with $\widehat{u}_{n+1} = 0$

**Algorithm 1**

$\kappa \leftarrow (\epsilon \cdot 1/2 \cdot \max_{i \in N_{1/2}} u_i)/n$
$z^* \leftarrow 0; t \leftarrow \emptyset$
**for** $j = h+1, \dots n$ **do**
  **if** $\widehat{z}(h,j) < q_j \cdot u_j$ **then**
    **if** $q_j \cdot u_j > z^*$ **then**
      $z^* \leftarrow q_j \cdot u_j$
      $t \leftarrow \{j\}$
    **end if**
  **else**
    **if** $\widehat{z}(h,j) > z^*$ **then**
      $z^* \leftarrow \widehat{z}(h,j)$
      $t \leftarrow t(\widehat{P}(h, \widehat{C}_{i,j}, \widehat{L}_{i,j}))$
    **end if**
  **end if**
**end for**
**return** $t$

and $q_{n+1} = 1$. Then, for $i \in [n]$ and $j > i$ the scaled DP problem is defined as

$$\widehat{z}_\kappa(i,j) = \max_{C \in [\overline{C}(i)], L \in [G]} (\kappa \cdot C + u_j) \cdot \widehat{P}(i,C,L) \cdot q_j. \tag{2.9}$$

where $\overline{C}(i) = \sum_{i' \in [i]} u_{i'}$. Let

$$\widehat{C}_{i,j}, \widehat{L}_{i,j} = \operatorname*{arg\,max}_{C \in [\overline{C}(i)], L \in [G]} (\kappa \cdot C + u_j) \cdot \widehat{P}(i,C,L) \cdot q_j.$$

Note that $t(\widehat{P}(i, \widehat{C}_{i,j}, \widehat{L}_{i,j}) \cup \{j\})$ returns an optimal test by replacing $u_i$ with $\widehat{u}_i$ and adding the constraint that for any $\ell \in [i+1, \dots, n] \setminus \{j\}$, $\ell$ is not pooled into the test, while $j$ is pooled into it. From Lemma 3, we know that it suffices to evaluate $\widehat{z}(i,j)$ for $i \in [h]$ and $j \in \{h+1, \dots, n\}$ in order to evaluate $\widehat{z}_\kappa(n, n+1)$ as at most one individual from $[n] \setminus N_{1/2}$ may be pooled into the test.

The following lemma establishes an upper bound on a value of $\kappa$ that suffices to bound the relative error of solutions of $\widehat{z}_\kappa$ in approximating the optimal test within a given $\epsilon > 0$.

**Lemma 5.** *Let $t^*$ be an optimal test with $\sum_{\ell \in t^*} u_\ell = C^*$ and $|t^*| = L^*$. For a given $\epsilon > 0$, there exist $i \in [h]$, $j \in \{h, \dots, n+1\}$, with $i < j$, and for $\overline{t} = t^* \setminus \{j\}$ and $\kappa \leqslant \frac{\epsilon \max_{i \in \overline{t}} q_i \cdot u_i}{n}$ such that*

$$\widehat{z}_\kappa(i,j) \geqslant (1-\epsilon) \cdot C^* \cdot P(n, C^*, L^*).$$

*Proof.* First note

$$\sum_{i \in t^*} u_i - \kappa \sum_{i \in t^*} \widehat{u}_i = \sum_{i \in t^*} u_i - \kappa \sum_{i \in t^*} \lfloor u_i/\kappa \rfloor \leqslant \sum_{i \in t^*} u_i - \kappa \sum_{i \in t^*} (u_i/\kappa - 1) \leqslant \kappa n$$

where the last inequality follows since $|t^*| \leqslant n$.

Let $j$ be the individual with the smallest probability of being healthy in $t^*$ by breaking ties with respect to individuals that have higher index and let $i$ be the individual with the highest index in $t^* \setminus \{j\}$. We denote with $\widehat{t} \subseteq [i]$ the set that maximizes Equation (2.9). Then, we get that

$$\begin{aligned}\widehat{z}_\kappa(i,j) = q_j \left(\kappa \sum_{\ell \in \widehat{t} \setminus \{j\}} \widehat{u}_\ell + u_j\right) q_{\widehat{t} \setminus \{j\}} &= q_j \left(\kappa \sum_{\ell \in t^* \setminus \{j\}} \widehat{u}_\ell + u_j\right) q_{t^* \setminus \{j\}} \\ &\geqslant \left(1 - \frac{n\kappa}{\sum_{i \in t^*} u_i}\right) \cdot q_{t^*} \sum_{i \in t^*} u_i,\end{aligned}$$

where the first inequality follows from optimality of $\widehat{t}$ under the scaled utilities. Thus, to ensure an $\epsilon$-approximate solution, we need

$$\frac{n\kappa}{\sum_{\ell \in t^*} u_\ell} \leqslant \epsilon \Leftrightarrow \kappa \leqslant \frac{\epsilon \cdot \sum_{\ell \in t^*} u_\ell}{n}.$$

Thus, it suffices to choose

$$\kappa \leqslant \frac{\epsilon \cdot \max_{\ell \in [t^* \setminus \{j\}]} u_\ell \cdot q_\ell}{n} \leqslant \frac{\epsilon \cdot q_{t^*} \cdot \sum_{i \in t^*} u_i}{n} \leqslant \frac{\epsilon \cdot \sum_{\ell \in t^*} u_i}{n}$$

□

Since $t^*$ is not known, we should choose a value for $\kappa$ that satisfies the above lemma. Note that due to Lemma 3, we know that $t^* \cap ([n] \setminus N_{1/2}) \leqslant 1$.

Algorithm 1, which is an FPTAS for the optimal test, due to Lemma 3, fixes an individual $j \in [n+1] \setminus N_{1/2}$ that is pooled into the optimal test, where $j = n+1$ indicates the case that $t^* \cap ([n] \setminus N_{1/2}) = 0$. Hence, we can apply Lemma 4 by setting $i = h$ and thus $\overline{t} \subseteq N_{1/2}$. Since for each $\ell \in N_{1/2}$, $q_\ell \geqslant 1/2$, we have that $\max_{\ell \in \overline{t}} q_\ell \cdot u_\ell \geqslant 1/2 \max_{\ell \in \overline{t}} \cdot u_\ell$. Thus, we can choose $\kappa$ such that

$$\kappa = \frac{\epsilon \cdot 1/2 \cdot \max_{\ell \in [\overline{t}]} u_\ell}{n} \leqslant \frac{\epsilon \cdot \max_{\ell \in [\overline{t}]} q_\ell \cdot u_\ell}{n} \leqslant \frac{\epsilon \cdot q_{t^*} \cdot \sum_{i \in t^*} u_i}{n}$$

where the last inequality follows from optimality of $t^*$.

Now, we show that Algorithm 1 is an FPTAS for the optimal test.

First note that if $|t^*| = 1$, then Algorithm 1 finds the optimal test in Lines 5-7. Hence, we focus on the case that $|t^*| > 1$. using Lemma 3, we distinguish into two cases:

**Case I:** $|t^* \setminus N_{1/2}| = 0$**.** For each given $\epsilon > 0$, $\kappa$ satisfies the supposition of Lemma 5. So following Lemma 5 with $\overline{C} = \overline{C}(h) = \sum_{\ell \in N_{1/2}} \geqslant \sum_{\ell \in t^*} \widehat{u}_\ell$, for $\widehat{C}_{h,n+1}$ and $\widehat{L}_{h,n+1}$, we have

$$\begin{aligned}u(t) = \widehat{z}_\kappa(h, n+1) = \kappa \cdot \widehat{C}_{h,n+1} \cdot \widehat{P}(h, \widehat{C}_{h,n+1}, \widehat{L}_{h,n+1}) &\geqslant (1-\epsilon) \cdot C^* \cdot P(h, C^*, L^*) \\ &\geqslant (1-\epsilon) \cdot C^* \cdot P(n, C^*, L^*)\end{aligned}$$

**Case II:** $|t^* \setminus N_{1/2}| = 1$**.** Then, for each $\epsilon > 0$, the choice of $\kappa$ for some $j \in [n] \setminus N_{1/2}$ satisfies

$$\widehat{z}_\kappa(h, j) \geqslant (1 - \epsilon) \cdot C^* \cdot P(n, C^*, L^*)$$

where the inequality follows form Lemma 5. The algorithm must determine $j$ since it enumerates all elements of $[n] \setminus N_{1/2}$. in the main loop.

The complexity of the algorithm is determined by at most $n$ evaluations of Equation (2.9). Hence,

$$O\left(n^3\overline{C}\right) \subseteq O\left(n^3 \sum_{\ell \in N_{1/2}} \frac{u_\ell}{\kappa}\right) \subseteq O\left(\frac{n^5}{\epsilon}\right).$$

□

Our goal is to approximate the optimal non-overlapping testing regime when there are $B$ available tests. Given the FPTAS of Lemma 4 for the single test case, a natural greedy approach is the following: design a test in each step by applying the FPTAS for the single test case to the remaining individuals. In other words, in each iteration we greedily find the test that approximates the optimal test over the available individuals using the FPTAS, add this test to the testing regime, disregard all individuals that are included in this test and continue to create greedy tests in the same fashion for the remaining individuals until the budget is exhausted. The above procedure results in an non-overlapping testing regime with at most $B$ tests, as we never consider individuals that have already been included in a test. We refer this algorithm as $\epsilon$-*Greedy*, where $\epsilon$ is the error tolerance used in the FPTAS algorithm[3]. We show that this algorithm gives a $5/(1 - \epsilon)$ approximation to the optimal non-overlapping testing regime.

**Theorem 2.** *$\epsilon$-Greedy returns a $5/(1 - \epsilon)$-approximate non-overlapping testing regime.*

*Proof.* Let $T$ be the testing regime that is returned by $\epsilon$-Greedy and $T^*$ be an optimal non-overlapping testing regime. Without loss of generality, let $N' = \{1, \ldots, n'\}$ be the set of individuals that are pooled into a test in $T$. Then,

$$u(T^*) = \sum_{j \in [B]} u(t_j^*) = \sum_{j \in [B]} q_{t_j^*} \cdot \left(\sum_{i \in N'} I_{i,j}^{t^*} \cdot u_i\right) + \sum_{j \in [B]} q_{t_j^*} \cdot \left(\sum_{i \in [n] \setminus N'} I_{i,j}^{t^*} \cdot u_i\right),$$

and

$$u(T) = \sum_{j \in [B]} u(t_j) = \sum_{j \in [B]} q_{t_j} \cdot \left(\sum_{i \in N'} I_{i,j}^{T} \cdot u_i\right) = \sum_{j \in [B]} q_{t_j} \cdot \left(\sum_{i \in t_j} u_i\right).$$

Now, let $T'$ be a testing regime such that $t'_j = t_j^* \setminus (t_j^* \cap N')$. In other words, $T'$ is created by removing from $T^*$ any individual in $N'$. Notice that since every $t'_j$ consists of individuals that are not included in any test in $T$, it means that all the individuals in $t'_j$ are available at the $j$-th iteration of the

[3]When $G$ is constant, we can efficiently compute optimal $t^*$ at each step of the approach above via brute force. We call this algorithm *Greedy*.

algorithm, and thus we get that for each $j \in [B]$, $u(t_j) \geqslant (1-\epsilon)u(t'_j)$, as otherwise the algorithm would have chosen $t'_j$ instead of $t_j$, at the $j$-th iteration. Thus, we get that

$$u(T) = \sum_{j\in[B]} u(t_j) \geqslant (1-\epsilon) \sum_{j\in[B]} u(t'_j) = (1-\epsilon) \cdot u(T'),$$

and also note that

$$\begin{aligned} u(T') = \sum_{j\in[B]} q_{t'_j} \cdot \left( \sum_{i\in[n]\setminus N'} I^{T'}_{i,j} \cdot u_i \right) &= \sum_{j\in[B]} q_{t'_j} \cdot \left( \sum_{i\in[n]\setminus N'} I^{T^*}_{i,j} \cdot u_i \right) \\ &\geqslant \sum_{j\in[B]} q_{t^*_j} \cdot \left( \sum_{i\in[n]\setminus N'} I^{T^*}_{i,j} \cdot u_i \right) \end{aligned}$$

where the second equality follows from the fact that $I^{T'}_{i,j} = I^{T^*}_{i,j}$ for any $i \in [n] \setminus N'$ and $j \in [B]$, and the last inequality follows from the fact that $q_{t'_j} \geqslant q_{t^*_j}$ since $t'_j \subseteq t^*_j$ for any $j \in [B]$. Thus,

$$u(T) \geqslant (1-\epsilon) \cdot \sum_{j\in[B]} q_{t^*_j} \cdot \left( \sum_{i\in[n]\setminus N'} I^{T^*}_{i,j} \cdot u_i \right).$$

From all the above we have

$$\begin{aligned} \frac{u(T^*)}{u(T)} &= \frac{\sum_{j\in[B]} q_{t^*_j} \cdot \left(\sum_{i\in N'} I^{T^*}_{i,j} \cdot u_i\right) + \sum_{j\in[B]} q_{t^*_j} \cdot \left(\sum_{i\in[n]\setminus N'} I^{T^*}_{i,j} \cdot u_i\right)}{u(T)} \\ &\leqslant \frac{\sum_{j\in[B]} q_{t^*_j} \cdot \left(\sum_{i\in N'} I^{T^*}_{i,j} \cdot u_i\right) + \frac{u(T)}{1-\epsilon}}{u(T)} \\ &= \frac{\sum_{j\in[B]} q_{t^*_j} \cdot \left(\sum_{i\in t^*_j\cap N'} u_i\right)}{u(T)} + \frac{1}{1-\epsilon} \\ &\leqslant \frac{\sum_{j\in[B]} \left(\sum_{i\in t^*_j\cap N'} q_i \cdot u_i\right)}{u(T)} + \frac{1}{1-\epsilon} \\ &= \frac{\sum_{i\in N'} q_i \cdot u_i}{u(T)} + \frac{1}{1-\epsilon} \end{aligned} \tag{2.10}$$

where the second inequality follows since $q_i \geqslant q_{t_j}$ when $i$ is included in $t_j$.

In what follows, we will show that each test, $t_j \in T$ obtains at least a $\frac{(1-\epsilon)}{4}$ ratio of the maximal possible utility to be gained from individuals in $t_j$. In other words, we show that the following holds: $u(t_j) = q_{t_j} \cdot \sum_{i\in t_j} u_i > \frac{(1-\epsilon)}{4} \sum_{i\in t_j} q_i u_i$. To do so, we first consider the case where there exists $i' \in t_j$ such that $q_{i'} < 1/2$. From the definition of the greedy algorithm, we know that

$$(1-\epsilon) \cdot q_{t_j\setminus\{i'\}} \cdot \sum_{i\in t_j\setminus\{i'\}} u_i \leqslant q_{t_j} \cdot \sum_{i\in t_j} u_i$$

otherwise the algorithm would return $t_j \setminus \{i'\}$ instead of $t_j$ at step $j$ and also,

$$(1-\epsilon) \cdot q_{i'} \cdot u_{i'} \leqslant q_{t_j} \cdot \sum_{i \in t_j} u_i$$

as otherwise the algorithm would return $\{i'\}$ instead of $t_j$ at the $j$-th iteration. Moreover, from Lemma 3, we know that $q_{t_j \setminus \{i'\}} \geqslant 1/2$ since $q_{i'} < 1/2$. Thus, we get that

$$\begin{aligned} q_{t_j} \cdot \sum_{i \in t_j} u_i \geqslant \frac{(1-\epsilon)}{2}(q_{t_j \setminus \{i'\}} \cdot \sum_{i \in t_j \setminus \{i\}} u_i + q_{i'} \cdot u_{i'}) &\geqslant \frac{(1-\epsilon)}{2}\left(\frac{1}{2} \cdot \sum_{i \in t_j \setminus \{i\}} u_i + q_{i'} \cdot u_{i'}\right) \\ &= \frac{(1-\epsilon)}{4}\left(\sum_{i \in t_j \setminus \{i'\}} u_i + q_{i'} \cdot u_{i'}\right) \\ &\geqslant \frac{(1-\epsilon)}{4} \sum_{i \in t_j} q_i u_i. \end{aligned}$$

As a second case, assume that for any $i \in t_j$, $q_i \geqslant 1/2$. We show that for each $i \in t_j$, $q_{t_j \setminus \{i\}} \geqslant 1/4$. If $q_{t_j} \geqslant 1/2$, then indeed $q_{t_j \setminus \{i\}} \geqslant 1/4$. Otherwise, we do the following: In a set $S$ we add individuals that are included in $t_j$, except for $i$, one at a time until $q_S \geqslant 1/2$ and $q_{S \cup \{i\}} < 1/2$ (as $q_{t_j} < 1/2$ notice that such an $S$ should exist). Then, from Lemma 3, we get that $q_{t_j \setminus \{S \cup i\}} \geqslant 1/2$, and hence $q_{t_j \setminus \{i\}} = q_S \cdot q_{t_j \setminus \{S \cup i\}} \geqslant 1/4$. Thus, we have that

$$q_{t_j} \cdot \sum_{i \in t_j} u_i = \sum_{i \in t_j} q_{t_j \setminus \{i\}} \cdot q_i \cdot u_i \geqslant \frac{1}{4} \cdot \sum_{i \in t_j} q_i \cdot u_i \geqslant \frac{(1-\epsilon)}{4} \sum_{i \in t_j} q_i \cdot u_i.$$

We see that in either case, $q_{t_j} \cdot \sum_{i \in t_j} u_i > \frac{(1-\epsilon)}{4} \cdot \sum_{i \in t_j} q_i \cdot u_i$ hence we get:

$$u(T) = \sum_{j \in [B]} q_{t_j} \cdot \left(\sum_{i \in t_j} u_i\right) > \sum_{j \in [B]} \left(\frac{(1-\epsilon)}{4} \cdot \sum_{i \in t_j} q_i \cdot u_i\right) = \frac{(1-\epsilon)}{4} \sum_{i \in N'} q_i u_i$$

Along with Equation (2.10), we get that $u(T^*)/u(T) \leqslant 5/(1-\epsilon)$ and the theorem follows. □

Note that one can combine Theorem 1 and Theorem 2 to see that $\epsilon$-Greedy gives a constant-factor approximation to the optimal *overlapping* testing regime.

## 2.4 Practical Algorithmic Implementations

As discussed above, even the problem of allocating a single test is computationally hard. In order to make the computation of testing regimes tractable for our pilot study, we formulate the problem of non-overlapping testing as optimization problems that we solve using commercial solvers. The problem of allocating a single test can be formulated as a mixed-integer conic optimization program (MICP), and solved using a commercial conic solver. This implementation is used by our Greedy algorithm in our

simulations and our pilot study. When multiple tests are to be allocated, we formulate a mixed-integer linear program (MILP) that approximates an optimal non-overlapping solution and can be solved by any MILP solver. The MILP formulation approximates existing exponential constraints with piecewise-linear functions that can be formulated as a collection of mixed integer linear constraints. The accuracy of this approximation can be adjusted by tuning the number $K$ of segments of the piecewise-linear function, at the cost of introducing more (integer) variables and thus the time to solve the program. We provide practical (additive) approximation guarantees for this *Approx algorithm* depending on parameter $K$.

**The optimization program.** We state the non-linear program for determining optimal non-overlapping testing regimes and describe a conic formulation for the single test case as well as a mixed-integer formulation for approximating optimal non-overlapping with one or more tests. We can assume that the testing budget $B$ is at most the population size $n$, and so pool sizes lies between 1 and $G$. For each test $j \in [B]$, we introduce an indicator vector $x^j \in \{0,1\}^n$ with $x_i^j = 1$ if individual $i$ is included in $j$ and $x_i^j = 0$ otherwise, and let variable $w^j$ denote its expected utility $w^j = u \cdot x^j \prod_{i\in[n]} q_i^{x_i^j}$. We impose pool sizes between 1 and $G$ with constraints $1 \leqslant \sum_{i\in[n]} x_i^j \leqslant G$ for all $j \in [B]$, and non-overlapping testing with constraints $\sum_{j\in[B]} x_i^j \leqslant G$ for all $i \in [n]$. Our objective is to maximize welfare $\sum_{j\in[B]} w^j$. In order to isolate the non-linear elements of the optimization problem, we reformulate the convex program with additional variables below; variables $l^j$ denote the log of $w^j$, and variables $y^j$ and $z^j$ allow us to isolate the non-linear elements of the expressions into constraints (2.11b) and (2.11d).

$$\max \quad \sum_{j\in[B]} w^j \tag{2.11a}$$

$$\text{s.t.} \quad w^j = \exp l^j, \qquad \forall j \in [B], \tag{2.11b}$$

$$l^j = y^j + \sum_{i\in[n]} x_i^j \log q_i, \quad \forall j \in [B], \tag{2.11c}$$

$$y^j = \log z^j, \qquad \forall j \in [B], \tag{2.11d}$$

$$z^j = u \cdot x^j, \qquad \forall j \in [B], \tag{2.11e}$$

$$\sum_{j\in[B]} x_i^j \leqslant 1, \qquad \forall i \in [n], \tag{2.11f}$$

$$1 \leqslant \sum_{i\in[n]} x_i^j \leqslant G, \qquad \forall j \in [B], \tag{2.11g}$$

$$x_i^j \in \{0,1\}, \qquad i \in [n], \forall j \in [B] \tag{2.11h}$$

## A Conic Program for a Single Test

Suppose we wish to allocate a single test. In this setting, we can eliminate the exponential constraint (2.11b) by changing the objective to $\max l^1$. The remaining non-linear constraints (2.11d) can be relaxed to $y^j \leqslant \log z^j$ without affecting the outcome, and formulated as conic constraints $(z^j, 1, y^j) \in K_{\exp}$, where $K_{\exp}$ is the exponential cone defined as $K_{\exp} = \{(x_1, x_2, x_3) \mid x_1 \geqslant x_2 e^{x_3/x_2}, x_2 > 0\} \cup \{(x_1, 0, x_3) \mid x_1 \geqslant 0, x_3 \leqslant 0\}$. The resulting mixed-integer conic optimization program can be

solved efficiently[4] with conic solvers such as MOSEK (`https://mosek.com`). In our practical implementation, the Greedy algorithm repeatedly solves a conic program to allocate a single test.

## A Mixed-Integer Linear Programming Approximation

If we wish to allocate more than one test, the problem no longer admits a conic formulation. Instead, we formulate a mixed-integer linear program (MILP) that finds an approximately optimal non-overlapping solution. In order to make the problem tractable, we assume that the utility vector $u$ is integral and non-negative. This assumption is benign, as the problem is invariant to scaling of utilities. We describe how the non-linear constraints (2.11b) and (2.11d) can respectively be captured exactly and approximately by a collection of integer linear constraints.

**Handling the logarithmic constraints.** We can replace (2.11d) with integer linear constraints as follows. Fix some test $j \in [B]$. Note that $z^j$ takes integral values in the range $[L, U]$, where $L = \min_i u_i$ and $U = G \max_i u_i$. We introduce an indicator vector $\gamma^j \in \{0,1\}^{[L,U]}$ indexed by $k \in [L, U]$ with constraints $\sum_{k\in[L,U]} \gamma^j_k = 1$ and $\sum_{k\in[L,U]} k \cdot \gamma^j_k = z^j$ to encode which value $z$ holds, and ensure $y^j = \log(z^j)$ with the constraint $y^j = \sum_{k\in[L,U]} \log(k) \cdot \gamma^j_k$.

**Approximating the exponential constraints.** We now describe how to approximate (2.11b) from above by a piecewise-linear function $f$ using integer linear constraints. Fix some test $j \in [B]$. Note first that we can relax the equality in (2.11b) to $w^j \leqslant \exp(l^j)$ without affecting the outcome. The variable $l^j$ takes values between $A = \min_i(\log u_i) + G \min_i(\log q_i)$ and $B = \log(G \max_i u_i) + \max_i(\log q_i)$ (and these values will be generically non-integral). We approximate $\exp$ from above by a piecewise-linear function $f : [A, B] \to \mathbb{R}$ with $K$ linear segments. (Here the parameter $K$ is given exogenously.) Partitioning $[A, B]$ into $K$ parts $[c_k, c_{k+1}], k \in [K]$, we define the $k$-th line segment as the linear function $f_k(x) = a_k x + b_k$ on domain $[c_k, c_{k+1}]$ with slope $a_k = \frac{\exp c_{k+1} - \exp c_k}{c_{k+1} - c_k}$ and residual $b_k = \exp c_{k+1} - a_k c_{k+1}$. Note that the number of integer variables in the MILP increases with $K$, so this parameter must be chosen judiciously. Moreover, given a fixed number of segments $K$, we wish to determine a partitioning of $[A, B]$ that minimizes the approximation error $\varepsilon = \max_{x\in[A,B]}(f(x) - \exp(x))$. In our implementation, we apply binary search techniques to numerically determine the partition of $[A, B]$ such that the error $\max_{x\in[c_k - c_{k+1}]}(f_k(x) - \exp(x))$ is the same for all parts $[c_k, c_{k+1}]$, which minimizes $\varepsilon$. We introduce indicator vectors $\delta^j \in \{0,1\}^K$ to encode in which part $[c_k, c_{k+1}]$ the value of $l^j$ lies, as well as the vector $v^j \in \mathbb{R}^K$ whose $k$-th entry agrees with $l^j$ if $l^j$ lies in the $k$-th part, and is $0$ otherwise. This is guaranteed by constraints $\sum_{k\in[K]} \delta^j_k = 1$, $l^j = \sum_{k\in[K]} v^j_k$ and $c_k \cdot \delta^j_k \leqslant v^j_k \leqslant c_{k+1} \cdot \delta^j_k, \forall k \in [K]$. Finally, we require that $w^j \leqslant f_k(l^j)$ for the $k$-th part $[c_k, c_{k+1}]$ that $l^j$ lies in. This is expressed by constraint $w^j \leqslant \sum_{k\in[K]} a_k v^j_k + b_k \cdot \delta^j_k$.

**Bounding the approximation error.** Recall that the piecewise-function $f$ with $K$ segments approximates $\exp$ on domain $[A, B]$ from above with error $\varepsilon$. Let $\sigma(x) = \sum_{j\in[B]} \exp(l^j)$ and $\sigma'(x) = \sum_{j\in[B]} f(l^j)$ respectively denote the corresponding objective values of the convex program (2.11) and the MILP described above for testing $x$. Let $x^*$ denote an optimal non-overlapping testing, so $x^*$

[4]Example running times are shown in Table 2.1 and Table 2.3.

maximizes $\sigma$, and $x'$ be an optimal solution for the MILP. Clearly, $x^*$ and $x'$ are both feasible for both programs and satisfy $\sigma(x') \leqslant \sigma(x^*)$ as well as $\sigma'(x^*) \leqslant \sigma'(x')$. By construction of $f$, we have $\sigma(x) \leqslant \sigma'(x)$ and $\sigma(x) \geqslant \sigma'(x) - \varepsilon B$, which implies $\sigma(x^*) \leqslant \sigma'(x^*) \leqslant \sigma'(x') \leqslant \sigma(x') + \varepsilon B$. Here $\varepsilon$ is the additive approximation error of $f$ with regard to $\exp$. Hence, $0 \leqslant \sigma(x^*) - \sigma(x') \leqslant \varepsilon B$. This allows us to compute a bound on the additive gap between the welfare achieved by the optimal solution of our MILP and the optimal non-overlapping testing.

## 2.5 Empirical Results

In this section, we evaluate the accuracy and running times of the Greedy and Approx algorithms on populations that reflect real-world scenarios. .

In our first experiments, we run both algorithms on the real data from the pilot study that executed at IPICYT in September 2020 by Lock et al. [2022] with budgets $B \in \{2, 4, \ldots, 12\}$. [5] As the health probabilities observed during the pilot study were high, this illustrates the efficacy of Greedy in times of lower disease incidence. In order to study how well Greedy performs when faced with higher infection rates, we also ran experiments on synthetic data in which health probabilities range from $0.5$ to $1$ and budgets $B \in \{2, 4, \ldots, 12\}$. Moreover, while our partnering testing laboratory performed pooled testing with saliva samples, which has a pool size limit of $G = 5$, we also study outcomes when pool sizes are limited to $G = 10$ (the limit for nasopharyngeal swabs).

In our synthetic experiments, we showcase the average-case behavior of Greedy and Approx by generating random populations of size $n = 150$. Health probabilities are drawn independently and uniformly at random from the interval $[0.5, 1]$, and utilities are drawn from a normal distribution that was fitted to the utilities observed in our pilot study. We then run Greedy and Approx on each population for each pool size $G \in \{5, 10\}$ and for all testing budgets $B \in \{2, 4, \ldots, 10\}$, recording the welfare achieved for both algorithms, as well as their running times (in milliseconds).[6]

For all experiments, we record the true welfares achieved by the testing regimes returned by both algorithms, and not the objective values of the underlying MILP and conic optimization problems, as the latter will be an approximation of the true welfare. For Approx, we tune the parameter $K$ of the MILP formulation so that the additive approximation guarantee (cf. Section 2.4) is small ($K = 25$ for the experiments on pilot data, and $K = 20$ for the experiments on synthetic data).

**Results.** Table 2.1 lists the welfares achieved by Approx and Greedy on our pilot study data for pool size constraint $G = 5$, as well as the running times for both algorithms and the approximation guarantee achieved by Approx. Table 2.2 shows analogous results for pool size constraint $G = 10$. We observe that Greedy achieves near-optimal welfare for budgets up to $10$ for both pool sizes. Moreover, while the running time of Approx appears to increase exponentially with the testing budget $B$, Greedy scales linearly and is extremely fast (even for larger populations and testing budgets). This makes Greedy

[5] As our mixed-integer program formulation is designed to admit integral utilities only, and the problem of computing testing regimes is invariant to scaling utilities, we first scale up the utilities of all individuals in the population by a factor of $50$, and then round the resulting number to the nearest integer. Choosing a larger scaling factor increases the running time, as the number of variables in the MILP increases (cf. Section 2.4).

[6] The experiments were run on an AWS EC2 instance type 'c6g.8xlarge' with 32 vCPUs and 64GiB memory. Gurobi 9.5.0 was used for the MILP, and MOSEK 10 for the MICP.

| | MILP | | | Greedy | | |
|---|---|---|---|---|---|---|
| Budget | Welfare | Guarantee | Time | Welfare | Apx To Optimal | Time |
| 2 | 461.24 | 0.18 | 292 ms | 461.22 | 1.000028 | 49 ms |
| 6 | 1292.04 | 0.53 | 1911 ms | 1291.91 | 1.000097 | 48 ms |
| 10 | 2070.82 | 0.89 | 5345 ms | 2070.58 | 1.000115 | 71 ms |
| 14 | 2814.82 | 1.25 | 2550656 ms | 2814.50 | 1.000115 | 400 ms |
| 18 | 3524.78 | 1.61 | 1523318 ms | 3524.24 | 1.000154 | 2718 ms |
| 22 | 4189.57 | 1.97 | 569366 ms | 4188.85 | 1.000170 | 42869 ms |
| 26 | 4790.46 | 2.32 | 12160741ms | 4789.37 | 1.000227 | 747631 ms |
| 30 | 4805.24 | 2.68 | 69541821 ms | 4789.37 | 1.003313 | 750206 ms |
| 34 | 4816.37 | 3.04 | 408803099 ms | 4789.37 | 1.005637 | 757912 ms |

Table 2.1: Summary showing welfare and computation time for Approx and Greedy on the pilot data with a population of $n = 130$ and pool size constraint $G = 5$, with testing budgets $B \in \{2, 6, \ldots, 34\}$. We also state the additive approximation guarantee of MILP (compared to optimal non-overlapping welfare).

attractive for implementations that rely on a quick turnaround, run on 'budget hardware' or wish to avoid costly cloud computing services.

In our synthetic experiments with lower health probabilities, Greedy performs as well as Approx when $G = 5$, and remains competitive also when $G = 10$. Figures 2.1 and 2.2 plot the mean welfares achieved by both algorithms for pool size constraints $G = 5$ and $G = 10$, as well as the welfare ratios. For the latter, we divide the welfare achieved by Approx by the welfare of Greedy for each population, and depict the resulting ratios as black dots. Tables 2.3 and 2.4 list the mean welfares and running times of both algorithms, as well as the approximation guarantee of Approx, for $G \in \{5, 10\}$.

Comparing our experiments with different pool sizes $G \in \{5, 10\}$, we see that increasing pool sizes from 5 to 10 significantly increases mean welfare if health probabilities are very high (cf. Tables 2.1 and 2.2). This effect is less pronounced in our experiment with synthetic data, in which participants have lower health probabilities on average (cf. Tables 2.3 and 2.4). These results suggest that the pool size limit of $5$ imposed by saliva sampling, as opposed to the limit of $10$ for nasopharyngeal samples, may be considered a limitation in some scenarios, and institutions may wish to weigh the positives and negatives of saliva and nasopharyngeal sampling carefully.

Here we show figures and summary tables for our experiments comparing Approx and Greedy on the pilot study data with pool size constraint $G = 10$, and on synthetic populations of size $n = 200$ and pool size constraints $G \in \{5, 10\}$.

## 2.6 Discussion

This work introduces a novel utility-based approach to pooled testing in resource-constrained environments. In this setting, we provide strong theoretical and empirical performance guarantees that further

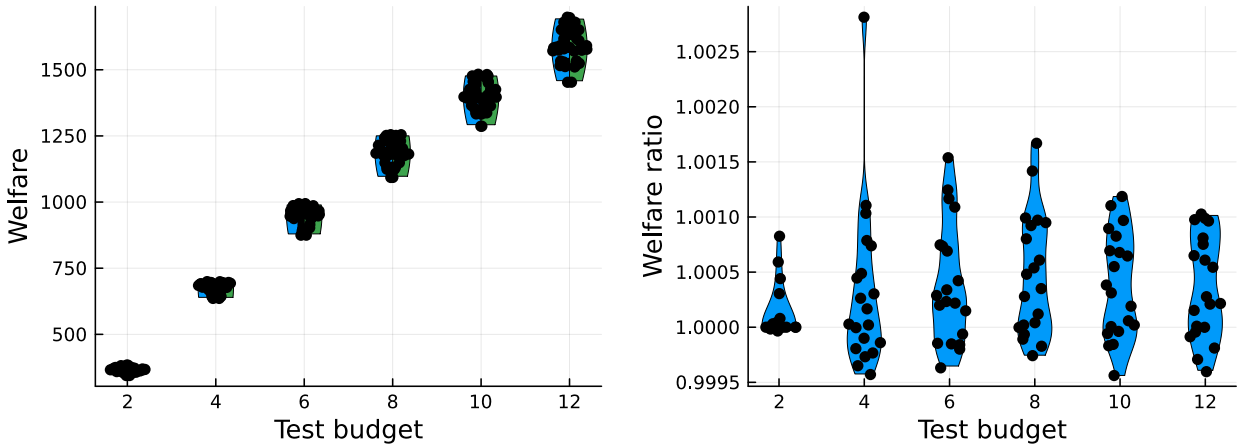


Figure 2.1: Outcomes of Greedy and Approx on synthetic data with $n = 200$, pool size bound $G = 5$ and testing budgets $B \in \{2, 4, \ldots, 12\}$. Left: Welfares achieved by MILP (left regions, blue) and Greedy (right regions, red). Right: Ratios between the welfares of Approx and Greedy. In both figures, each black dot corresponds to one of the 20 randomly generated populations.

| | MILP | | | Greedy | | |
|---|---|---|---|---|---|---|
| Budget | Welfare | Guarantee | Time | Welfare | Apx To Optimal | Time |
| 2 | 866.58 | 0.47 | 470 ms | 866.47 | 1.000127 | 52 ms |
| 6 | 2391.49 | 1.42 | 3504 ms | 2390.99 | 1.000210 | 50 ms |
| 10 | 3775.35 | 2.36 | 1781766 ms | 3774.20 | 1.000305 | 87 ms |
| 14 | 4696.30 | 3.31 | 19293052 ms | 4678.52 | 1.003800 | 279 ms |
| 18 | 4741.30 | 4.25 | 62784039 ms | 4678.52 | 1.013417 | 672 ms |
| 22 | 4770.00 | 5.19 | 88422845 ms | 4678.52 | 1.019554 | 1086 ms |
| 26 | 4790.29 | 6.13 | 133956040 ms | 4678.52 | 1.023889 | 289 ms |
| 30 | 4805.25 | 7.08 | 284989519 ms | 4678.52 | 1.027088 | 1215 ms |
| 34 | 4816.15 | 8.03 | 545021914 ms | 4678.52 | 1.029417 | 1175 ms |

Table 2.2: Summary showing welfare and computation time for the MILP and Greedy on pilot study data (with a population of $n = 130$) and pool size constraint $G = 10$ with testing budgets $B \in \{2, 6, \ldots, 34\}$. We also state the additive approximation guarantee of MILP (compared to optimal non-overlapping welfare).

justify the implementation of non-overlapping testing regimes beyond their essential logistical simplicity. We evaluate our algorithms Greedy and Approx on real date and notice that Greedy performs almost optimally and is significantly faster than our alternative MILP implementation.

There are many directions for future work. On a theoretical level, there is a gap between our upper bound of 4 and lower bound of 7/6 on the gain of overlaps, and an upper bound of 5 on the approximation factor of the greedy algorithm. Beyond the tight approximation bound of $e$ for Greedy that we establish

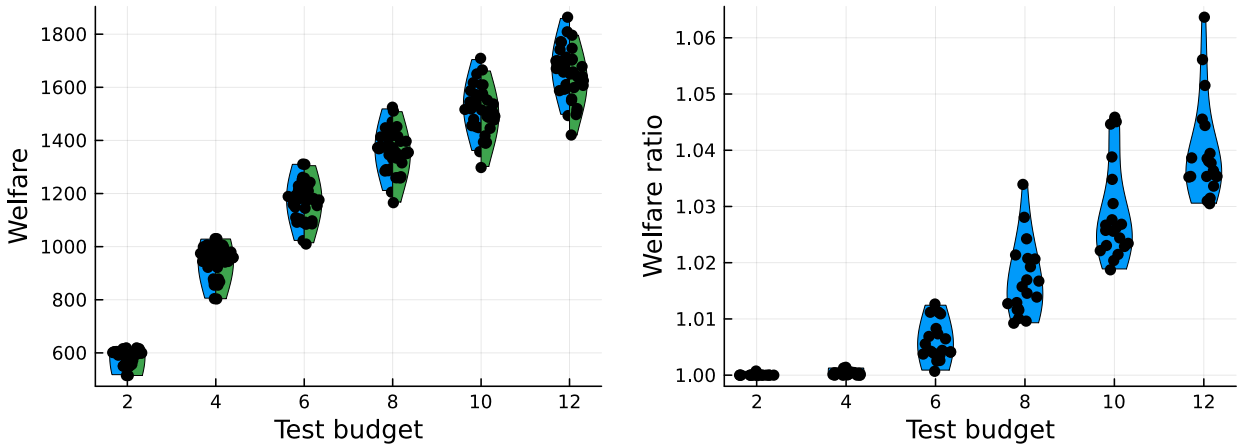


Figure 2.2: Outcomes of Greedy and Approx on synthetic data with $n = 200$, pool size bound $G = 10$ and testing budgets $B \in \{2, 4, \ldots, 12\}$. Left: Welfares achieved by Approx (left regions, blue) and Greedy (right regions, red). Right: Ratios between the welfares of Approx and Greedy. In both figures, each black dot corresponds to one of the 20 randomly generated populations.

for the case with identical utilities for individuals, we expect that tighter bounds are achievable when utilities take a fixed number of values (e.g. for dichotomous or trichotomous populations). On a more practical level, the overall testing and re-integration policy we propose is static in nature, as we consider the one-shot setting where a testing budget is to be fully utilized by a policymaker. But testing can be dynamic, with allocations chosen adaptively as a function of previous test results, and it is valuable to understand what potential benefits this extended functionality can bring. Additionally, policymakers potentially have access to different types of tests, each with different associated costs and performance (i.e., pool size and sensitivity), and providing optimal budget-constrained allocations in this heterogeneous test setting is a key open question.

Most importantly, although most countries have eased COVID-19 restrictions, we hope that the valuable insight in performance and efficacy of our welfare-maximizing testing regimes can help better protect resource-constrained communities during future outbreaks of infectious diseases.

| | Approx | | | Greedy | |
|---|---|---|---|---|---|
| Budget | Welfare | Guarantee | Time | Welfare | Time |
| 2 | 366.90 | 0.55662 | 285 ms | 366.85 | 35 ms |
| 4 | 675.51 | 1.11324 | 837 ms | 675.29 | 67 ms |
| 6 | 947.86 | 1.66986 | 2585 ms | 947.48 | 98 ms |
| 8 | 1188.80 | 2.22648 | 9933 ms | 1188.24 | 134 ms |
| 10 | 1399.83 | 2.7831 | 202447 ms | 1399.30 | 175 ms |
| 12 | 1584.99 | 3.33972 | 949195 ms | 1584.43 | 226 ms |

Table 2.3: Experiment summary on synthetic data with pool size bound $G = 5$ and testing budgets $B \in \{2, 4, \ldots, 12\}$. Welfares and times are averaged over 20 randomly generated populations. We also state the additive approximation guarantee of Approx (compared to optimal non-overlapping welfare).

| | Approx | | | Greedy | |
|---|---|---|---|---|---|
| Budget | Welfare | Guarantee | Time | Welfare | Time |
| 2 | 587.20 | 1.23777 | 389 ms | 587.16 | 33 ms |
| 4 | 941.65 | 2.47554 | 2089 ms | 941.34 | 83 ms |
| 6 | 1177.26 | 3.71331 | 9297 ms | 1170.46 | 171 ms |
| 8 | 1368.17 | 4.95108 | 33184 ms | 1345.93 | 265 ms |
| 10 | 1533.75 | 6.18884 | 75437 ms | 1491.34 | 379 ms |
| 12 | 1680.95 | 7.42661 | 350039 ms | 1617.25 | 566 ms |

Table 2.4: Experiment summary on synthetic data with pool size bound $G = 10$ and testing budgets $B \in \{2, 4, \ldots, 12\}$. Welfares and times are averaged over 20 randomly generated populations. We also state the additive approximation guarantee of Approx (compared to optimal non-overlapping welfare).

# Chapter 3

# Grading Schemes and Students' Performance in Consecutive Tests

## 3.1 Introduction

Student evaluations and grading play an integral and influential role in every individual's academic experience. Naturally, there has been widespread debate among researchers and policy-makers about the efficacy of various grading systems such as *letter v.s. number grades*. For instance, coarse-grained grading schemes (i.e., letter grades) are considered to be less noisy indicators of performance, and stronger signals of status, and consequently, are the norm in North American universities. At the same time, there is also growing awareness that the grade itself affects performance independent of student ability, i.e., the grades are "not just an output of the educational process, they may also be an input" [Gray and Bunte, 2022]. For example, empirical evidence suggests the disclosure of midterm grades may motivate or demotivate students to perform better in a future exam, controlling for other effects. In light of this evidence, it is clear that the design of a grading system must be a deliberate choice that takes into account student welfare in addition to other extraneous factors [Guskey, 2011]. In this work, we take an analytical approach and study the design of an optimal grading system with a particular focus on numeric v.s. uniform letter grades.[1] As far as we are aware, this work is among the first to look at the problem of designing a grading scheme with the explicit objective of improving student performance in future tests. Our model captures the impact of grades on future performance via two well-motivated effects:

1. **Anchoring:** In any given test, students anchor themselves to a specific score or performance level based on their intrinsic ability which directly impacts their performance. We refer to this anchor as the intrinsic quality.

2. **(De)Motivation**: When the student's actual score falls above (below) their intrinsic quality, they get (de)motivated and subsequently, their expectation increases (decreases) for future tests. This is a phenomenon that has been widely noticed in practice [Deci et al., 1999, Dev, 1997, Cameron and Pierce, 1994].

[1] We use the term uniform letter grades to refer to letter grading schemes where each letter grade corresponds to an equal sized score range, e.g., $[90, 100] \rightarrow$A+, $[80, 90] \rightarrow$A-, and so on.

In this regard, our work departs from other papers in this area, where students are often modelled as status-maximizers [Dubey and Geanakoplos, 2010], i.e., their intrinsic motivation for a better grade stems from a desire to rank above their fellow students. Our model does not induce any artificial scarcity (status) and instead the fundamental friction is a result of noisy performance and how the same grading rule affects different students differently.

To better illustrate how different grading schemes impact student performance under our model, consider the case of two students with the same intrinsic quality $q_1 = q_2 = 85$. Due to random factors, the first student's score in the midterm is given by $s_1 = 81$ while the second student matches expectations and scores $s_2 = 85$. In this case, disclosing the numeric score may demotivate student 1, leading to an effective intrinsic quality for the final exam that is lower than $85$. On the other hand, under a coarser scheme, both students could receive a letter grade (say) $A-$ capturing all scores in the range $[80, 90]$, which limits the adverse effect on future performance. At the same time, a third student who's intrinsic quality is $q_3 = 91$ and whose midterm score is $s_3 = 89$ may also be bracketed into the same letter grade $A-$, leading to severe demotivation. In this scenario, the disclosure of the numeric grade would inform the third student that their performance was actually close to their intrinsic quality.

This example crucially illustrates that the way that students perceive a non-numerical grade plays a key role, and this often depends on how such grades are perceived in the outside world. We study a scheme of mapping letter grades to percentages, which is widely used in practice (see, e.g., University of Western Ontario [2022], Victoria University of Wellington [2022]), where the grade of the student is given by the midpoint of the interval containing all the scores that map to that letter grade. For example, if a student receives $A-$, which captures all scores in the range $[80, 90]$, the student effectively receives a grade of $85$, and this grade is what the student compares to her intrinsic quality.

Building on the ideas presented in this example, we develop a framework to compare various grading systems in an environment with *sequential testing*. This includes evaluations within a course, e.g., a midterm followed by a final exam, but also grading across related courses, e.g., a student taking Calculus 101 followed by Calculus 102. Since a student's intrinsic quality increases after a test if the grade received is higher than her intrinsic quality and decreases otherwise, our aim is

> ...to compare different grading schemes and choose the one that provides a higher quality improvement (or a lower quality degradation).

**Our results.** In this work, we compare the numerical scoring scheme, where the student learns her exact score in an evaluation, to uniform letter grading schemes, where the interval of scores is partitioned into $T$ equal-length intervals mapping to different letter grades (and each interval is represented by its midpoint). Note that under uniform (or even non-uniform) letter grading schemes, one cannot simply assign all students a grade of $100$ to maximally motivate them: for example, if the entire range of $[0, 100]$ is mapped to one letter grade ($T = 1$), all students would receive a grade of $50$ due to the midpoint representation; e.g., as mentioned before, it prevents the instructor from simply assigning a grade of $100$ to all the students regardless of their scores.

First, we theoretically study the case where two sequential evaluations take place. We show that under natural conditions, numerical scoring and all uniform letter grading schemes have equal performance when the motivational and demotivational effects are equally strong, and otherwise, either numerical scoring outperforms all uniform letter grading schemes or the opposite happens. By assuming additional conditions, such as when the intrinsic qualities of the students follow a uniform distribution, we can

conclude that numerical scoring outperforms all uniform letter grading schemes when the demotivational effect is stronger than the motivational effect, and the opposite happens when the demotivational effect is weaker than the motivational effect. Since there is significant evidence that negative events have a greater impact than positive events [Baumeister et al., 2001, Coleman et al., 1987], we expect the demotivational effect to be stronger than the motivational effect; thus, our results are in favour of numerical scoring.

Next, we empirically compare numerical scoring to uniform letter grading schemes. Under two sequential evaluations, we observe that numerical scoring continues to outperform uniform letter grading when the demotivational effect is stronger (and the opposite continues to hold when the motivational effect is stronger), even under more realistic conditions than in our theoretical analysis, such as when the true qualities of the students follow a (truncated) normal distribution. However, surprisingly, when more than two evaluations take place, the effect is reversed. Even after just six sequential evaluations, uniform letter grading begins to outperform numerical scoring when the demotivational effect is stronger (and the opposite holds when the motivational effect is stronger). In the intermediate stage between these two regimes, there is another surprising effect: with four sequential evaluations, numerical scoring outperforms uniform letter grading regardless of which effect is stronger!

Our results indicate that the choice of the grading scheme depends on the application at hand: with fewer evaluations (e.g., courses with just a few tests or shorter education programs with just a few semesters), numerical scoring may be better, while with many evaluations (e.g., courses with weekly tests or longer education programs), uniform letter grading may be better. At a high level, although our work draws on literature from fields such as Economics and Psychology, it provides a fundamental perspective on the question of student grading within the framework of multi-agent systems, i.e., where each student is modeled as an agent whose behavior depends on the decisions made by the system. Our results open up the possibility of designing grading systems that are easy to implement, approximately-optimal, and take into account students' incentives.

**Related work.** There is a rich literature on comparing grading schemes using various objectives. However, to the best of our knowledge, none of these papers study the objective of improving student quality that we focus on.

Several works have studied, both theoretically and empirically, how the effort exerted by students for an evaluation depends on the grading scheme to be used [Paredes, 2017, Brownback, 2018, Main and Ost, 2014, Czibor et al., 2020]. For example, when using pass/fail grading, a student may try hard enough to pass (with high probability), but not any harder. Our work is orthogonal to this: we focus on effect of the outcome of one evaluation on the student motivation in *subsequent* evaluations.

Another related work is that of Sikora [2015], who also compares grading schemes, but his goal is to study the tradeoff between conveying the most information about the student's true quality and minimizing noise due to factors unrelated to the true quality, not the (de)motivational effects of the grading scheme in subsequent evaluations. In our work, the task of keeping the grades "consistent" with the actual performance is indirectly performed by our use of the midpoint representation.

Rohe et al. [2006] and Bloodgood et al. [2020] also study how the grading scheme used may impact students' psychological well-being and stress levels, but do not focus on the impact of this in subsequent evaluations.

## 3.2 Model

Define $[k] = \{1, \ldots, k\}$ for $k \in \mathbb{N}$. We introduce a model in which the grading scheme used in one evaluation can motivate or demotivate students, affecting their performance in future evaluations.

**True qualities.** A student begins with an intrinsic (true) quality $q$ drawn from a (nonatomic) prior $\mathcal{Q}$ with probability density function (PDF) $f_{\mathcal{Q}}(\cdot)$. For simplicity, let the support of $\mathcal{Q}$ be $[0, 1]$.

**Scores.** There is a *score model* $\mathcal{S}$ such that the numerical performance (score) of a student with true quality $q$ in the first evaluation, denoted $s \in [0, 1]$, is drawn from the (nonatomic) distribution $\mathcal{S}(q)$ with PDF $f_{\mathcal{S}}(\cdot; q)$. We focus on score models in which the expected score of a student is equal to their true quality, i.e., $\mathsf{E}_{s\sim\mathcal{S}(q)}[s] = q$ for all $q \in [0, 1]$.

**Grades.** A grading scheme is a function $B : [0, 1] \to [0, 1]$ that maps the score to a grade.

*Letter grading.* A *letter grading* scheme $B_{\vec{c}}$ is specified by a vector $\vec{c} = (c_0 = 0, c_1, \ldots, c_{T-1}, c_T = 1)$, for some $T \in \mathbb{N}$ (referred to as the number of grades) and $c_i \geqslant c_{i-1}$ for all $i \in [T]$, and is given by $B_{\vec{c}}(s) = \frac{c_{i-1}+c_i}{2}$ for all $i \in [T]$ and $s \in [c_{i-1}, c_i)$. That is, it partitions $[0, 1)$ into finitely many disjoint intervals (one for each grade) and maps a score to the midpoint of the interval containing it.

*Uniform letter grading.* We are particularly interested in the *uniform letter grading* (ULG) scheme. For a given number of grades $T \in \mathbb{N}$, uniform letter grading with $T$ grades, denoted $\mathrm{ULG}_T$, is specified by $c_i = i/T$ for each $i \in [T]$. In other words, it partitions $[0, 1)$ into $T$ equal-length intervals. We will use $\Delta(T) = 1/T$ to denote the length of the interval, dropping $T$ from the argument when it is clear from the context. Formally, we have that for all $s \in [0, 1)$,[2]

$$\mathrm{ULG}_T(s) = (\lfloor s/\Delta \rfloor + 1/2) \cdot \Delta.$$

For instance, $\mathrm{ULG}_{10}$ maps all scores in $[0, 0.1)$ to $0.05$, all scores in $[0.1, 0.2)$ to $0.15$, and so on. We restrict our focus to uniform grading schemes for two reasons: a) it is straightforward and easy to implement in practice; b) given that different institutions following different grading schemes, this allows us to broadly compare letter and number grading without getting lost in the minutiae. Further our assumption that each letter grade maps to the midpoint of an interval is common practice across universities [University of Western Ontario, 2022, Victoria University of Wellington, 2022] as well as the literature [McEwan et al., 2021, Nisbet, 1975]. More generally, it is consistent with the practice of assigning a score or grade-point to each letter grade.

*Numerical scoring.* We will compare (uniform) letter grading to *numerical scoring* (NS), given by $\mathrm{NS}(s) = s$ for all $s \in [0, 1]$. Under numerical scoring, scores are not rounded to any grades. This can also be viewed as the limit of uniform letter grading with $T \to \infty$ grades.

**(De)motivation.** The grades affect students' level of motivation in subsequent evaluations. Under grading scheme $B$, a student compares their true quality $q$ to the obtained grade $B(s)$. If the grade is higher than the true quality, the student experiences a motivational boost, but in the converse case, gets demotivated. We model this by assuming that the effective true quality of the student for the next evaluation changes to $q' = q + h(q, B(s))$, where

$$h(q, B(s)) = \begin{cases} \alpha_m \cdot (B(s) - q), & \text{if } B(s) \geqslant q, \\ -\alpha_d \cdot (q - B(s)), & \text{if } B(s) < q. \end{cases}$$

[2]Because we assume nonatomic distributions, it does not matter what $\mathrm{ULG}_T(1)$ is. We will use the convention that $\mathrm{ULG}_T(1) = 1$.

We refer to $\alpha_m, \alpha_d \in \mathbb{R}_{\geqslant 0}$ as *motivation and demotivation coefficients*, respectively. Note that the amount of (de)motivation is proportional to the difference between the obtained grade and the true quality. In the next evaluation, the student obtains a score $s'$ drawn from $\mathcal{S}(q')$. We remark that when $\alpha_m, \alpha_d \in [0, 1]$, we automatically have $q' \in [0, 1]$; thus, we focus on this range of parameters.[3] Our choice of a linear model for demotivation follows from studies showing that student performance is linearly dependent on both external [Christensen and Menzel, 1998] and internal stimuli [Latham and Locke, 2007]. Additionally, even when the actual behaviour is more complex, our model serves as a first-order approximation when $(B(s) - q)$ is small.

**Goal.** Intuitively, we are interested in choosing grading schemes that achieve a higher increase (or a lower decrease) in the average student quality. Thus, we define the *performance* of a grading scheme $B$ as:

$$\mathrm{perf}(B) \triangleq \mathsf{E}_{q\sim\mathcal{Q}, s\sim\mathcal{S}(q)}[q' - q]$$

where $q' = q + h(q, B(s))$. Due to linearity of expectation,

$$\mathrm{perf}(B) = \mathsf{E}_{q\sim\mathcal{Q}, s\sim\mathcal{S}(q)}[q' - q] = \mathsf{E}_{q\sim\mathcal{Q}, s\sim\mathcal{S}(q)}[h(q, B(s))].$$

Thus, we compare $\mathsf{E}_{q\sim\mathcal{Q}, s\sim\mathcal{S}(q)}[h(q, B(s))]$ under numerical scoring and uniform letter grading. Hereinafter, we omit $q \sim \mathcal{Q}$ and $s \sim \mathcal{S}(q)$ from an expression of expectation, whenever it is clear from the context.

Note that for our theoretical analysis, we focus on the case of two evaluations. Later, we empirically study the case of more than two evaluations.

## 3.3 Uniform Letter Grading vs Numerical Scoring

In this section, we derive theoretical results for the performance of uniform letter grading schemes and numerical scoring, when students participate in two sequential evaluations and identify conditions under which numerical scoring outperforms every uniform letter grading scheme, and conditions under which the converse holds. Let us begin by introducing two useful definitions.

**Definition 1** (Jointly Symmetric Distributions)**.** *We say that the true quality prior $\mathcal{Q}$ and the score model $\mathcal{S}$ are* jointly symmetric *if $f_{\mathcal{Q}}(q) \cdot f_{\mathcal{S}}(s; q) = f_{\mathcal{Q}}(1 - q) \cdot f_{\mathcal{S}}(1 - s; 1 - q)$ for all $s, q \in [0, 1]$.*

Joint symmetry requires that true qualities and scores are symmetric across $[0, 1]$. That is, the probability of having true quality $q$ and receiving score $s$ should be the same as the probability of having true quality $1 - q$ and receiving score $1 - s$. If the true quality prior is uniform, then this means the score distribution $\mathcal{S}(q)$ should be the mirror image of the score distribution $\mathcal{S}(1 - q)$. Note that joint symmetry does not necessarily require symmetry of the "noise" contained in the score compared to the true quality. For example, we do not need $f_{\mathcal{S}}(s = 0.4; q = 0.5) = f_{\mathcal{S}}(s = 0.6; q = 0.5)$.

**Definition 2** (Symmetric Grading Scheme)**.** *We say that a grading scheme $B$ is* symmetric *if $B(1 - s) = 1 - B(s)$ for all $s \in [0, 1]$.*

[3] In principle, one can also use larger coefficients and truncate $q'$ to lie in $[0, 1]$.

The reader can check that numerical scoring (NS) and uniform letter grading schemes ($\mathrm{ULG}_T$ for any $T \in \mathbb{N}$) are symmetric. Our first result shows that under such symmetry, the performance of the grading scheme is linear in the difference between the motivation and demotivation coefficients. As we later show in Corollary 4, this allows us to compare numerical scoring to uniform letter grading.

**Theorem 3.** *When the true quality prior $\mathcal{Q}$ and the score model $\mathcal{S}$ are jointly symmetric, and the grading scheme $B$ is symmetric, then we have*

$$\mathrm{perf}(B) = \frac{\alpha_m - \alpha_d}{2} \cdot \mathsf{E}_{q\sim\mathcal{Q}, s\sim\mathcal{S}(q)}\Big[|q - B(s)|\Big]. \tag{3.1}$$

*Proof.* Note that due to $\mathcal{Q}$ and $\mathcal{S}$ being jointly symmetric, the pairs $(q, s)$ and $(1-q, 1-s)$ are sampled with equal density. Hence, we have that

$$\mathsf{E}\Big[h(q, B(s))\Big] = \frac{1}{2} \cdot \mathsf{E}\Big[h(q, B(s)) + h(1-q, B(1-s))\Big]. \tag{3.2}$$

Due to the symmetry of the grading scheme, we have $B(1-s) = 1 - B(s)$, which implies that the two terms $h(q, B(s))$ and $h(1-q, B(1-s))$ are motivation and demotivation by the same amount. Hence,

$$\begin{aligned}&\mathsf{E}\Big[h(q, B(s)) + h(1-q, B(1-s))\Big]\\ &= (\alpha_m - \alpha_d) \cdot \mathsf{E}\Big[|q - B(s)|\Big].\end{aligned}$$

Plugging this into Equation (3.2), we get the result. □

**Corollary 4.** *Assume that the true quality prior $\mathcal{Q}$ and the score model $\mathcal{S}$ are jointly symmetric. Then, all symmetric grading schemes have equal performance when $\alpha_m = \alpha_d$. Further, when $\alpha_m \neq \alpha_d$, for every $T \in \mathbb{N}$ one of the following conditions holds.*

1. *Uniform letter grading with $T$ grades is at least as good as numerical scoring when $\alpha_m > \alpha_d$, and the converse holds when $\alpha_m < \alpha_d$.*
2. *Uniform letter grading with $T$ grades is at least as good as numerical scoring when $\alpha_m < \alpha_d$, and the converse holds when $\alpha_m > \alpha_d$.*

*Proof.* The first claim regarding $\alpha_m = \alpha_d$ follows immediately from Equation (3.1). For the second claim regarding $\alpha_m \neq \alpha_d$, note that the comparison between numerical scoring and uniform letter grading with $T$ buckets reduces to the sign of $\mathsf{E}[|q - \mathrm{NS}(s)| - |q - \mathrm{ULG}_T(s)|]$, and depending on this sign, one of the two statements in the corollary holds. □

Corollary 4 tells us that having equal motivation and demotivation coefficients ($\alpha_m = \alpha_d$) is the turning point: between uniform letter grading with a fixed number of grades and numerical scoring, one is better when $\alpha_m < \alpha_d$ but the other becomes better when $\alpha_m > \alpha_d$. But it does not tell us *which* one is better in each case.

Our next result identifies a sufficient condition under which this dilemma is settled: uniform letter grading is better when $\alpha_m > \alpha_d$ and numerical scoring is better when $\alpha_m < \alpha_d$. To introduce this sufficient condition, we need to define the following natural property of the score model.

**Definition 3** (Ex-Ante Single-Peaked Score Model)**.** *We say that the score model $\mathcal{S}$ is* ex-ante single-peaked *if, for every $q \in [0,1]$, $f_{\mathcal{S}}(\cdot; q)$ is single-peaked with the peak at $q$, i.e., $f_{\mathcal{S}}(s; q) \leqslant f_{\mathcal{S}}(s'; q)$ for all $s \leqslant s' \leqslant q$ and $s \geqslant s' \geqslant q$.*

Intuitively, in an ex-ante single-peaked score model, scores closer to the true quality are more likely than scores farther from the true quality.

For a fixed $T$, we also denote with $\mathcal{D}$ the set of all pairs of true qualities and scores that belong to the same letter grade interval, i.e., $\mathcal{D} = \{(q, s) : \mathrm{ULG}_T(q) = \mathrm{ULG}_T(s)\}$. For example, if $T = 10$, $(q = 0.51, s = 0.59) \in \mathcal{D}$ but $(q = 0.51, s' = 0.49) \notin \mathcal{D}$.

**Theorem 5.** *Fix any $T \in \mathbb{N}$. Assume that the true quality prior $\mathcal{Q}$ and the score model $\mathcal{S}$ satisfy the following.*

1. *$\mathcal{Q}$ and $\mathcal{S}$ are jointly symmetric;*
2. *$\mathcal{S}$ is ex-ante single-peaked; and*
3. $\mathsf{E}\Big[|q - s| \;\Big|\; (q, s) \in \mathcal{D}\Big]$
   $\leqslant \mathsf{E}\Big[|q - \mathrm{ULG}_T(s)| \;\Big|\; (q, s) \in \mathcal{D}\Big]$.

*Then, the first implication of Corollary 4 holds. That is, uniform letter grading with $T$ grades is at least as good as numerical scoring when $\alpha_m > \alpha_d$, the converse holds when $\alpha_m < \alpha_d$, and the two have equal performance when $\alpha_m = \alpha_d$.*

Before diving into the proof, let us make a remark regarding the third technical condition in Theorem 5. The technical condition states that, averaged over all such pairs, the true quality is closer to the score than to the midpoint of the interval that they both belong to. Later, we show that this condition is satisfied in two natural cases. Intuitively, if the score distribution is sufficiently concentrated near the true quality, the expected distance between the score and the true quality will be sufficiently small, satisfying the condition. Let us now turn to the proof of Theorem 5.

*Proof.* Given Theorem 3, we simply need to show that $\mathsf{E}\Big[|q - s|\Big] \leqslant \mathsf{E}\Big[|q - \mathrm{ULG}_T(s)|\Big]$. We already assume that this holds conditioned on $(q, s) \in \mathcal{D}$. Hence, we only need to show that it also holds conditioned on $(q, s) \notin \mathcal{D}$. We show this given the additional single-peakedness property.

We show that, conditioned on $(q, s) \notin \mathcal{D}$, the desired equation actually holds for every $q \in [0, 1]$, and, thus, in expectation over $q \sim \mathcal{Q}$ too. Fix any $q \in [0, 1]$. Note that

$$\begin{aligned}
&\mathsf{E}\Big[|q - s| \;\Big|\; (q, s) \notin \mathcal{D}\Big] \\
&\quad = \Pr\Big[\mathrm{ULG}_T(s) < \mathrm{ULG}_T(q) \;\Big|\; (q, s) \notin \mathcal{D}\Big] \\
&\qquad \cdot \mathsf{E}\Big[q - s \;\Big|\; \mathrm{ULG}_T(s) < \mathrm{ULG}_T(q)\Big] \\
&\qquad + \Pr\Big[\mathrm{ULG}_T(s) > \mathrm{ULG}_T(q) \;\Big|\; (q, s) \notin \mathcal{D}\Big] \\
&\qquad \cdot \mathsf{E}\Big[s - q \;\Big|\; \mathrm{ULG}_T(s) > \mathrm{ULG}_T(q)\Big]
\end{aligned}$$

$$
\begin{aligned}
&\leqslant \Pr\Big[\mathrm{ULG}_T(s) < \mathrm{ULG}_T(q) \;\Big|\; (q,s) \notin \mathcal{D}\Big] \\
&\qquad \cdot \mathsf{E}\Big[q - \mathrm{ULG}_T(s) \;\Big|\; \mathrm{ULG}_T(s) < \mathrm{ULG}_T(q)\Big] \\
&\quad + \Pr\Big[\mathrm{ULG}_T(s) > \mathrm{ULG}_T(q) \;\Big|\; (q,s) \notin \mathcal{D}\Big] \\
&\qquad \cdot \mathsf{E}\Big[\mathrm{ULG}_T(s) - q \;\Big|\; \mathrm{ULG}_T(s) > \mathrm{ULG}_T(q)\Big] \\
&= \mathsf{E}\Big[|q - \mathrm{ULG}_T(s)| \;\Big|\; (q,s) \notin \mathcal{D}\Big],
\end{aligned}
$$

where the first transition holds because

$$
[0,1]^2 \setminus \mathcal{D} = \{(q,s) : \mathrm{ULG}_T(s) < \mathrm{ULG}_T(q)\} \cup \{(q,s) : \mathrm{ULG}_T(s) > \mathrm{ULG}_T(q)\}
$$

and the second transition holds due to linearity of expectation and because the single-peakedness assumption implies

$$
\begin{aligned}
&\mathsf{E}\Big[s \;\Big|\; \mathrm{ULG}_T(s) < \mathrm{ULG}_T(q)\Big] \\
&\qquad \geqslant \mathsf{E}\Big[\mathrm{ULG}_T(s) \;\Big|\; \mathrm{ULG}_T(s) < \mathrm{ULG}_T(q)\Big], \\
\text{(and)}\;\; &\mathsf{E}\Big[s \;\Big|\; \mathrm{ULG}_T(s) > \mathrm{ULG}_T(q)\Big] \\
&\qquad \leqslant \mathsf{E}\Big[\mathrm{ULG}_T(s) \;\Big|\; \mathrm{ULG}_T(s) > \mathrm{ULG}_T(q)\Big].
\end{aligned}
$$

This completes the proof. □

In Theorem 5, we argued that single-peakedness of $\mathcal{S}$ establishes the desired inequality of $\mathsf{E}\Big[|q - s|\Big] \leqslant \mathsf{E}\Big[|q - \mathrm{ULG}_T(s)|\Big]$ at least conditioned on $(q,s) \notin \mathcal{D}$, leaving only the case of $(q,s) \in \mathcal{D}$, which was stated as an assumption in. Next, we show that if the true quality prior $\mathcal{Q}$ is uniform over $[0,1]$, and it satisfies two natural assumptions then the desired inequality also holds conditioned on $(q,s) \in \mathcal{D}$.

**Definition 4** (Ex-Post Single-Peaked Score Model)**.** *We say that the score model $\mathcal{S}$ is* ex-post single-peaked *if, for every $s \in [0,1]$, $f_{\mathcal{S}}(s;\cdot)$ is single-peaked with the peak at $s$, i.e., $f_{\mathcal{S}}(s;q) \leqslant f_{\mathcal{S}}(s;q')$ for all $s \leqslant q' \leqslant q$ and $q \leqslant q' \leqslant s$.*

**Definition 5** (Probabilistic Single-Dipped Score Model)**.** *We say that the score model $\mathcal{S}$ is* probabilistic single-dipped *if, for every $x \in [0,1]$, $\Pr\Big[s \in [q,x] \cup [x,q] \;\Big|\; q\Big]$ (let us call this $p(x,q)$) is single-dipped in $q$ with the dip at $q = x$, i.e., $p(x,q) \leqslant p(x,q')$ for all $x \leqslant q' \leqslant q$ and $q \leqslant q' \leqslant x$.*

Before we state the next theorem, we further partition $\mathcal{D}$ into two sub-spaces, $\mathcal{D}^{\text{same}}$ and $\mathcal{D}^{\text{opp}}$, such that $\mathcal{D}^{\text{same}}$ contains the set of all pairs of true qualities and scores such that either both are at most or both are at least the midpoint of their common letter grade interval, i.e.

$$
\begin{aligned}
\mathcal{D}^{\text{same}} = \{(q,s) : {} & q, s \leqslant \mathrm{ULG}_T(q) = \mathrm{ULG}_T(s) \\
& \vee\, q, s \geqslant \mathrm{ULG}_T(q) = \mathrm{ULG}_T(s)\}
\end{aligned}
$$

and $\mathcal{D}^{\text{opp}} = \mathcal{D} \setminus \mathcal{D}^{\text{same}}$. For example, when $T = 10$, $(q = 0.54, s = 0.51) \in \mathcal{D}^{\text{same}}$, but $(q = 0.54, s' = 0.56) \in \mathcal{D}^{\text{opp}}$. We are now ready to state the result.

**Theorem 6.** *Fix arbitrary $T \in \mathbb{N}$. Assume the following regarding the true quality prior $\mathcal{Q}$ and the score model $\mathcal{S}$.*

1. *$\mathcal{Q}$ is uniform over $[0, 1]$;*
2. *$\mathcal{Q}$ and $\mathcal{S}$ are jointly symmetric;*
3. *$\mathcal{S}$ is ex-ante and ex-post single-peaked, and probabilistic single-dipped; and*
4. $\Pr\Big[(q, s) \in \mathcal{D}^{\text{same}}\Big] \geqslant 2(\gamma + 1) \cdot \Pr\Big[(q, s) \in \mathcal{D}^{\text{opp}}\Big]$, *where* $\gamma = \max_{a,b\in[0,1]} \frac{f_{\mathcal{S}}(a;b)}{f_{\mathcal{S}}(b;a)}$.

*Then, the first implication of Corollary 4 holds. That is, uniform letter grading with $T$ grades is at least as good as numerical scoring when $\alpha_m > \alpha_d$, the converse holds when $\alpha_m < \alpha_d$, and the two have equal performance when $\alpha_m = \alpha_d$.*

Let us now understand the assumptions in Theorem 6. A natural choice of $\mathcal{S}$ under which Assumptions 3 and 4 in Theorem 6 are satisfied is when $\mathcal{S}(q)$ is a symmetric distribution around $q$, i.e., the noise in the score follows a symmetric zero-mean distribution. Further, for such a score model, we have $\gamma = 1$, so Assumption 4 becomes $\Pr[(q, s) \in \mathcal{D}^{\text{same}}] \geqslant 4 \cdot \Pr[(q, s) \in \mathcal{D}^{\text{opp}}]$.

More general, from the definitions of $\mathcal{D}^{\text{same}}$ and $\mathcal{D}^{\text{opp}}$, when the variance of the score distribution is sufficiently small, we can expect $\Pr[(q, s) \in \mathcal{D}^{\text{same}}]$ to be much higher than $\Pr[(q, s) \in \mathcal{D}^{\text{opp}}]$. For further intuition, see Figure 3.1.

Before we prove Theorem 6, we state the integral version of the well-known Chebyshev's inequality and its two useful implications.

**Lemma 6** (Integral Chebyshev Inequality)**.** *If functions $f, g : [a, b] \to \mathbb{R}_{\geqslant 0}$ are either both non-increasing or both non-decreasing, then*

$$\frac{1}{b-a}\int_a^b f(x)g(x)\,\mathrm{d}x \geqslant \left(\frac{1}{b-a}\int_a^b f(x)\,\mathrm{d}x\right)\cdot\left(\frac{1}{b-a}\int_a^b g(x)\,\mathrm{d}x\right).$$

*If one of them is non-decreasing while the other is non-increasing, the inequality is reversed.*

The following inequality is obtained by substituting $f(x) = x$ (and thus, $\frac{1}{b-a}\int_a^b f(x)\,\mathrm{d}x = \frac{a+b}{2}$) into Lemma 6.

**Lemma 7.** *If $g : [a, b] \to \mathbb{R}_{\geqslant 0}$ is a non-increasing function, then we have*

$$\int_a^b xg(x)\,\mathrm{d}x \leqslant \frac{a+b}{2}\cdot\int_a^b g(x)\,\mathrm{d}x,$$

*and the inequality is reversed if $g$ is a non-decreasing function.*

If $g$ is a probability density function over $[a, b]$, then $\int_a^b g(x)\,\mathrm{d}x = 1$, yielding the following (quite natural) implication.

**Lemma 8.** *Let $X$ be a random variable over $[a, b]$ with a non-increasing probability density function $g : [a, b] \to \mathbb{R}_{\geqslant 0}$. Then, $\mathsf{E}[X] \leqslant (a + b)/2$, and the inequality is reversed if $g$ is non-decreasing.*

Finally, we use the following strengthening of the integral Chebyshev inequality when one of the functions is linear and the other is concave non-increasing.

**Lemma 9.** *Let $g : [a, b] \to \mathbb{R}_{\geqslant 0}$ be a concave function with $g(b) = 0$. Then, we have*

$$\int_a^b (b-x)g(x)\,\mathrm{d}x \leqslant \frac{2(b-a)}{3}\int_a^b g(x)\,\mathrm{d}x.$$

*Proof.* Due to concavity of $g$, we have

$$\int_x^b g(t)\,\mathrm{d}t \geqslant \frac{1}{2}(b-x)g(x).$$

Hence, we have

$$\begin{aligned}
\int_a^b \frac{1}{2}(b-x)g(x)\,\mathrm{d}x &\geqslant \int_{x=a}^b \int_{t=x}^b g(t)\,\mathrm{d}t\,\mathrm{d}x \\
&= \int_{t=a}^b \int_{x=a}^t g(t)\,\mathrm{d}x\,\mathrm{d}t && \text{(Fubini's theorem)} \\
&= \int_{t=a}^b (t-a)g(t)\,\mathrm{d}t \\
&= \int_{x=a}^b (x-a)g(x)\,\mathrm{d}x && \text{(Change of variable name)} \\
&= \int_a^b (b-a)g(x)\,\mathrm{d}x - \int_a^b (b-x)g(x)\,\mathrm{d}x.
\end{aligned}$$

Rearranging the terms yields the desired inequality. □

Now, we are ready to prove Theorem 6.

*Proof.* Given Theorem 5, we only need to show that

$$\begin{aligned}
&\mathsf{E}\Big[|q-s| - |q-\mathrm{ULG}_T(s)| \;\Big|\; (q,s) \in \mathcal{D}\Big] \\
&\quad = \Pr[(q,s) \in \mathcal{D}^{\mathrm{same}} \;\Big|\; (q,s) \in \mathcal{D}] \cdot \mathsf{E}\Big[|q-s| - |q-\mathrm{ULG}_T(s)| \;\Big|\; (q,s) \in \mathcal{D}^{\mathrm{same}}\Big] \\
&\qquad + \Pr[(q,s) \in \mathcal{D}^{\mathrm{opp}} \;\Big|\; (q,s) \in \mathcal{D}] \cdot \mathsf{E}\Big[|q-s| - |q-\mathrm{ULG}_T(s)| \;\Big|\; (q,s) \in \mathcal{D}^{\mathrm{opp}}\Big] \\
&\quad \leqslant 0.
\end{aligned} \tag{3.3}$$

Let us analyze the expected value of $|q-s| - |q-\mathrm{ULG}_T(s)|$ conditioned on both $(q,s) \in \mathcal{D}^{\mathrm{same}}$ and $(q,s) \in \mathcal{D}^{\mathrm{opp}}$ separately.

**Analyzing $\mathcal{D}^{\text{same}}$.** For $k \in \{0, 1, \ldots, T-1\}$, define $\ell(k) = k\Delta$, $m(k) = (k+1/2)\Delta$, and $h(k) = (k+1)\Delta$. These are respectively the lower end, midpoint, and upper end of the $k$-th grade interval under $\mathrm{ULG}_T$. Note that

$$\mathcal{D}^{\text{same}} = \Big\{(q,s) : (\ell(k) \leqslant q \leqslant s \leqslant m(k)) \vee (\ell(k) \leqslant s \leqslant q \leqslant m(k)) \vee \\ (m(k) \leqslant q \leqslant s < h(k)) \vee (m(k) \leqslant s \leqslant q < h(k)),\ k \in \{0, 1, \ldots, T-1\}\Big\}.$$

Fix an arbitrary $k \in \{0, 1, \ldots, T-1\}$; write $\ell$, $m$, and $h$ while omitting the fixed $k$ in the argument; and let us analyze the desired expression $|q-s| - |q - \mathrm{ULG}_T(s)|$ conditioned on each of the four cases for this fixed $k$ separately. We will derive bounds that will hold regardless of the value of $k$, and, therefore, also conditional on $(q,s) \in \mathcal{D}^{\text{same}}$ (i.e., aggregated across all $k$). Note that in each case, we have $\mathrm{ULG}_T(q) = \mathrm{ULG}_T(s) = m$.

1. $\ell \leqslant q \leqslant s \leqslant m$. In this case, $|q-s| - |q - \mathrm{ULG}_T(s)| = s - m$. Note that

$$\begin{aligned}
\mathsf{E}\Big[s - m \mid \ell \leqslant q \leqslant s \leqslant m\Big] &= \frac{\int_{q=\ell}^{m}\int_{s=q}^{m} f_{\mathcal{Q}}(q)\cdot f_{\mathcal{S}}(s;q)\cdot(s-m)\,\mathrm{d}s\,\mathrm{d}q}{\Pr[\ell \leqslant q \leqslant s \leqslant m]} \\
&= \frac{\int_{q=\ell}^{m} 1 \cdot \mathsf{E}\Big[s-m \;\Big|\; q, s \in [q,m]\Big]\cdot \Pr\Big[s \in [q,m] \;\Big|\; q\Big]\,\mathrm{d}q}{\int_{q=\ell}^{m} 1 \cdot \Pr\Big[s \in [q,m] \;\Big|\; q\Big]\,\mathrm{d}q} \\
&\leqslant \frac{-\int_{q=\ell}^{m}\left(\frac{m-q}{2}\right)\cdot \Pr\Big[s \in [q,m] \;\Big|\; q\Big]\,\mathrm{d}q}{\int_{q=\ell}^{m} 1 \cdot \Pr\Big[s \in [q,m] \;\Big|\; q\Big]\,\mathrm{d}q} \\
&\leqslant \frac{-\frac{1}{m-\ell}\cdot\left(\int_{q=\ell}^{m}\frac{m-q}{2}\,\mathrm{d}q\right)\cdot\left(\int_{q=\ell}^{m}\Pr\Big[s \in [q,m] \;\Big|\; q\Big]\,\mathrm{d}q\right)}{\int_{q=\ell}^{m} 1 \cdot \Pr\Big[s \in [q,m] \;\Big|\; q\Big]\,\mathrm{d}q} \\
&= -\frac{1}{(\Delta/2)}\int_{r=0}^{\frac{\Delta}{2}}\frac{r}{2}\,\mathrm{d}r = -\frac{\Delta}{8}.
\end{aligned}$$

   Here, the third transition holds because conditioned on a given value of $q$ and on $s \in [q,m]$, the distribution of $s \in [q,m]$ is single-peaked with peak at $q$ (Assumption 3). Hence, $\mathsf{E}[s|q, s \in [q,m]] \leqslant (q+m)/2$. The fourth transition is the integral Chebyshev inequality (Lemma 6), which holds because both $(m-q)/2$ and $\Pr\Big[s \in [q,m] \;\Big|\; q\Big]$ are non-negative, non-increasing functions of $q$ in $[\ell, m]$ (Assumption 3).

2. $\ell \leqslant s \leqslant q \leqslant m$. In this case, $|q-s| - |q - \mathrm{ULG}_T(s)| = 2q - m - s$. Note that

$$\begin{aligned}
&\mathsf{E}\Big[2q - m - s \mid \ell \leqslant s \leqslant q \leqslant m\Big] \\
&\quad= \frac{\int_{q=\ell}^{m}\int_{s=\ell}^{q} f_{\mathcal{Q}\times\mathcal{S}}(q,s)\cdot(2q-m-s)\,\mathrm{d}s\,\mathrm{d}q}{\Pr[\ell \leqslant s \leqslant q \leqslant m]} \\
&\quad= \frac{\int_{s=\ell}^{m} f_{\mathcal{S}}(s)\mathsf{E}\Big[2q-m-s \;\Big|\; s, q \in [s,m]\Big]\cdot\Pr\Big[q \in [s,m] \;\Big|\; s\Big]\,\mathrm{d}s}{\Pr[\ell \leqslant s \leqslant q \leqslant m]}.
\end{aligned}$$

Here, we use $f_{\mathcal{Q}\times\mathcal{S}}(q,s) = f_{\mathcal{Q}}(q)\cdot f_{\mathcal{S}}(s;q)$ to denote the joint probability density of $q$ and $s$, and $f_{\mathcal{S}}(s) = \int_{q=0}^{1} f_{\mathcal{Q}}(q) f_{\mathcal{S}}(s;q)\,\mathrm{d}q$ to denote the marginal probability density of $s$.

We argue that $\mathsf{E}\Big[2q - m - s \;\Big|\; s, q \in [s,m]\Big] \leqslant 0$. Intuitively, this is because the posterior distribution of $q \in [s,m]$ conditioned on a fixed value of $s$ and on $q \in [s,m]$ is single-peaked with peak at $s$ by Assumptions 1 and 3. Hence, $\mathsf{E}\Big[q \;\Big|\; s, q \in [s,m]\Big] \leqslant (s+m)/2$. Formally, this can be viewed as

$$\begin{aligned}\mathsf{E}\Big[2q - m - s \;\Big|\; s, q \in [s,m]\Big] &= \frac{\int_{q=s}^{m} f(q;s)(2q-m-s)\,\mathrm{d}q}{\int_{q=s}^{m} f(q;s)\,\mathrm{d}q} \\ &\leqslant \frac{\frac{1}{m-s}\cdot\left(\int_{q=s}^{m} f(q;s)\,\mathrm{d}q\right)\cdot\left(\int_{q=s}^{m}(2q-m-s)\,\mathrm{d}q\right)}{\int_{q=s}^{m} f(q;s)\,\mathrm{d}q} = 0,\end{aligned}$$

where $f(q;s) = \frac{f_{\mathcal{Q}}(q)\cdot f_{\mathcal{S}}(s;q)}{f_{\mathcal{S}}(s)}$ denotes the probability density of true quality being $q$ conditioned on the score being $s$; the second transition is the integral Chebyshev inequality (Lemma 6), which holds because $f(q;s)$ is a non-increasing function of $q$ whereas $2q-m-s$ is a non-decreasing function of $q$;[4,5] and the final transition holds because the second integral in the numerator is $0$.

3. $m \leqslant q \leqslant s < h$. In this case, $|q-s| - |q - \mathrm{ULG}_T(s)| = m + s - 2q$. Due to the same reasoning as in Case 2, we have that $\mathsf{E}\Big[m + s - 2q \;\Big|\; m \leqslant q \leqslant s < h\Big] \leqslant 0$.

4. $m \leqslant s \leqslant q < h$. In this case, $|q-s| - |q - \mathrm{ULG}_T(s)| = m - s$. Due to the same reasoning as in Case 1, we have that $\mathsf{E}\Big[m - s \;\Big|\; m \leqslant s \leqslant q < h\Big] \leqslant -\Delta/8$.

Let $p_1, p_2, p_3, p_4$ respectively denote the total probabilities of the above four cases across all values of $k \in \{0,1,\ldots,T-1\}$, conditioned on $(q,s) \in \mathcal{D}^{\mathrm{same}}$. Then, $p_1+p_2+p_3+p_4 = 1$. Because $f_{\mathcal{S}}(a;b)/f_{\mathcal{S}}(b;a) \leqslant \gamma$ for all $a, b \in [0,1]$, it follows that $p_1 \geqslant p_2/\gamma$ and $p_4 \geqslant p_3/\gamma$. Hence, $p_1 + p_4 \geqslant (p_2+p_3)/\gamma$. Using $p_1+p_2+p_3+p_4 = 1$, we get $p_1 + p_4 \geqslant 1/(\gamma+1)$.

Combining the analysis from the four cases above, we have

$$\mathsf{E}\Big[|q-s| - |q - \mathrm{ULG}_T(s)| \;\Big|\; (q,s) \in \mathcal{D}^{\mathrm{same}}\Big] \leqslant -(p_1+p_4)\cdot\frac{\Delta}{8} \leqslant -\frac{\Delta}{8(\gamma+1)}. \tag{3.4}$$

**Analyzing $\mathcal{D}^{\mathrm{opp}}$.** Note that

$$\mathcal{D}^{\mathrm{opp}} = \cup_{k\in\{0,1,\ldots,T-1\}} \left\{(q,s) : (\ell(k) \leqslant q \leqslant m(k) \leqslant s \leqslant h(k)) \vee (\ell(k) \leqslant s \leqslant m(k) \leqslant q \leqslant h(k))\right\}.$$

Fix an arbitrary $k \in \{0,1,\ldots,T-1\}$; as before, write $\ell$, $m$, and $h$ while omitting the fixed $k$ in the argument. Once again, we analyze the desired expression $|q-s| - |q - \mathrm{ULG}_T(s)|$ conditioned on each

[4]To see why $f(q;s) = \frac{f_{\mathcal{Q}}(q)\cdot f_{\mathcal{S}}(s;q)}{f_{\mathcal{S}}(s)}$ is non-increasing in $q$, note that the denominator does not depend on $q$ whereas the numerator is equal to $f_{\mathcal{S}}(s;q)$ (Assumption 1), which is non-increasing in $q$ (Assumption 3).

[5]Technically, integral Chebyshev inequality requires non-negative functions, and $2q-(m+s)$ can be negative when $q < (m+s)/2$. However, one can equivalently separate out the $-(m+s)$ term, apply the integral Chebyshev inequality to $2q$, and recombine with the $-(m+s)$ term to achieve the same conclusion.

of the two cases in the above expansion of $\mathcal{D}^{\text{opp}}$ for this fixed $k$ separately. We will derive bounds that will hold regardless of the value of $k$, and, therefore, also conditional on $(q, s) \in \mathcal{D}^{\text{opp}}$ (i.e., aggregated across all $k$). Note that we still have $\text{ULG}_T(q) = \text{ULG}_T(s) = m$.

1. $\ell \leqslant q \leqslant m \leqslant s \leqslant h$: In this case, we have $|q - s| - |q - \text{ULG}_T(s)| = s - m$. Note that

$$\mathsf{E}\Big[s - m \;\Big|\; \ell \leqslant q \leqslant m \leqslant s \leqslant h\Big] \leqslant \Delta/4. \tag{3.5}$$

   This is because $s \in [m, m + \Delta/2]$ and, due to single-peakedness of the score model and $q \leqslant m$, it is at most $m + \Delta/4$ in expectation.

2. $\ell \leqslant s \leqslant m \leqslant q \leqslant h$: In this case, we have $|q - s| - |q - \text{ULG}_T(s)| = m - s$, and the same reasoning as above shows that

$$\mathsf{E}\Big[m - s \;\Big|\; \ell \leqslant s \leqslant m \leqslant q \leqslant h\Big] \leqslant \Delta/4. \tag{3.6}$$

Combining Equations (3.10) and (3.11) and aggregating over all $k \in \{0, 1, \ldots, T - 1\}$, we get that

$$\mathsf{E}\Big[|q - s| - |q - \text{ULG}_T(s)| \;\Big|\; (q, s) \in \mathcal{D}^{\text{opp}}\Big] \leqslant \Delta/4. \tag{3.7}$$

Finally, combining Equations (3.9) and (3.12), we have that

$$\begin{aligned}
&\mathsf{E}\Big[|q - s| - |q - \text{ULG}_T(s)| \;\Big|\; (q, s) \in \mathcal{D}\Big] \\
&\quad \leqslant \Pr\Big[(q, s) \in \mathcal{D}^{\text{same}} \;\Big|\; (q, s) \in \mathcal{D}\Big] \cdot \left(-\frac{\Delta}{8(\gamma + 1)}\right) + \Pr\Big[(q, s) \in \mathcal{D}^{\text{opp}} \;\Big|\; (q, s) \in \mathcal{D}\Big] \cdot \frac{\Delta}{4} \leqslant 0,
\end{aligned}$$

where the final transition holds because $\Pr\Big[(q, s) \in \mathcal{D}^{\text{same}}\Big] \geqslant 2(\gamma + 1) \cdot \Pr\Big[(q, s) \in \mathcal{D}^{\text{opp}}\Big]$ (Assumption 4), which is equivalent to

$$\Pr\Big[(q, s) \in \mathcal{D}^{\text{same}} \;\Big|\; (q, s) \in \mathcal{D}\Big] \geqslant 2(\gamma + 1) \cdot \Pr\Big[(q, s) \in \mathcal{D}^{\text{same}} \;\Big|\; (q, s) \in \mathcal{D}\Big].$$

This completes the proof. □

*Proof of Theorem 6.* Given Theorem 5, we only need to show that

$$\begin{aligned}
&\mathsf{E}\Big[|q - s| - |q - \text{ULG}_T(s)| \;\Big|\; (q, s) \in \mathcal{D}\Big] \\
&\quad = \Pr[(q, s) \in \mathcal{D}^{\text{same}} \;\Big|\; (q, s) \in \mathcal{D}] \cdot \mathsf{E}\Big[|q - s| - |q - \text{ULG}_T(s)| \;\Big|\; (q, s) \in \mathcal{D}^{\text{same}}\Big] \\
&\qquad + \Pr[(q, s) \in \mathcal{D}^{\text{opp}} \;\Big|\; (q, s) \in \mathcal{D}] \cdot \mathsf{E}\Big[|q - s| - |q - \text{ULG}_T(s)| \;\Big|\; (q, s) \in \mathcal{D}^{\text{opp}}\Big] \\
&\quad \leqslant 0.
\end{aligned} \tag{3.8}$$

Let us analyze the expected value of $|q - s| - |q - \text{ULG}_T(s)|$ conditioned on both $(q, s) \in \mathcal{D}^{\text{same}}$ and $(q, s) \in \mathcal{D}^{\text{opp}}$ separately.

**Analyzing $\mathcal{D}^{\text{same}}$.** For $k \in \{0, 1, \ldots, T-1\}$, define $\ell(k) = k\Delta$, $m(k) = (k + 1/2)\Delta$, and $h(k) = (k+1)\Delta$. These are respectively the lower end, midpoint, and upper end of the $k$-th grade interval under $\text{ULG}_T$. Note that

$$\mathcal{D}^{\text{same}} = \Big\{(q,s) : (\ell(k) \leqslant q \leqslant s \leqslant m(k)) \vee (\ell(k) \leqslant s \leqslant q \leqslant m(k)) \vee \\ (m(k) \leqslant q \leqslant s < h(k)) \vee (m(k) \leqslant s \leqslant q < h(k)),\ k \in \{0, 1, \ldots, T-1\}\Big\}.$$

Fix an arbitrary $k \in \{0, 1, \ldots, T-1\}$; write $\ell$, $m$, and $h$ while omitting the fixed $k$ in the argument; and let us analyze the desired expression $|q - s| - |q - \text{ULG}_T(s)|$ conditioned on each of the four cases for this fixed $k$ separately. We will derive bounds that will hold regardless of the value of $k$, and, therefore, also conditional on $(q, s) \in \mathcal{D}^{\text{same}}$ (i.e., aggregated across all $k$). Note that in each case, we have $\text{ULG}_T(q) = \text{ULG}_T(s) = m$.

1. $\ell \leqslant q \leqslant s \leqslant m$. In this case, $|q - s| - |q - \text{ULG}_T(s)| = s - m$. Note that

$$\begin{aligned}
\mathsf{E}\Big[s - m \mid \ell \leqslant q \leqslant s \leqslant m\Big] &= \frac{\int_{q=\ell}^{m}\int_{s=q}^{m} f_{\mathcal{Q}}(q) \cdot f_{\mathcal{S}}(s;q) \cdot (s-m)\,\mathrm{d}s\,\mathrm{d}q}{\Pr[\ell \leqslant q \leqslant s \leqslant m]} \\
&= \frac{\int_{q=\ell}^{m} 1 \cdot \mathsf{E}\Big[s - m \;\Big|\; q, s \in [q,m]\Big] \cdot \Pr\Big[s \in [q,m] \;\Big|\; q\Big]\,\mathrm{d}q}{\int_{q=\ell}^{m} 1 \cdot \Pr\Big[s \in [q,m] \;\Big|\; q\Big]\,\mathrm{d}q} \\
&\leqslant \frac{-\int_{q=\ell}^{m} \left(\frac{m-q}{2}\right) \cdot \Pr\Big[s \in [q,m] \;\Big|\; q\Big]\,\mathrm{d}q}{\int_{q=\ell}^{m} 1 \cdot \Pr\Big[s \in [q,m] \;\Big|\; q\Big]\,\mathrm{d}q} \\
&\leqslant \frac{-\frac{1}{m-\ell} \cdot \left(\int_{q=\ell}^{m} \frac{m-q}{2}\,\mathrm{d}q\right) \cdot \left(\int_{q=\ell}^{m} \Pr\Big[s \in [q,m] \;\Big|\; q\Big]\,\mathrm{d}q\right)}{\int_{q=\ell}^{m} 1 \cdot \Pr\Big[s \in [q,m] \;\Big|\; q\Big]\,\mathrm{d}q} \\
&= -\frac{1}{(\Delta/2)} \int_{r=0}^{\frac{\Delta}{2}} \frac{r}{2}\,\mathrm{d}r = -\frac{\Delta}{8}.
\end{aligned}$$

   Here, the third transition holds because conditioned on a given value of $q$ and on $s \in [q, m]$, the distribution of $s \in [q, m]$ is single-peaked with peak at $q$ (Assumption 3). Hence, $\mathsf{E}[s|q, s \in [q, m]] \leqslant (q + m)/2$. The fourth transition is the integral Chebyshev inequality (Lemma 6), which holds because both $(m - q)/2$ and $\Pr\Big[s \in [q, m] \;\Big|\; q\Big]$ are non-negative, non-increasing functions of $q$ in $[\ell, m]$ (Assumption 3).

2. $\ell \leqslant s \leqslant q \leqslant m$. In this case, $|q - s| - |q - \text{ULG}_T(s)| = 2q - m - s$. Note that

$$\begin{aligned}
&\mathsf{E}\Big[2q - m - s \mid \ell \leqslant s \leqslant q \leqslant m\Big] \\
&\quad= \frac{\int_{q=\ell}^{m}\int_{s=\ell}^{q} f_{\mathcal{Q}\times\mathcal{S}}(q,s) \cdot (2q-m-s)\,\mathrm{d}s\,\mathrm{d}q}{\Pr[\ell \leqslant s \leqslant q \leqslant m]} \\
&\quad= \frac{\int_{s=\ell}^{m} f_{\mathcal{S}}(s)\mathsf{E}\Big[2q - m - s \;\Big|\; s, q \in [s,m]\Big] \cdot \Pr\Big[q \in [s,m] \;\Big|\; s\Big]\,\mathrm{d}s}{\Pr[\ell \leqslant s \leqslant q \leqslant m]}.
\end{aligned}$$

Here, we use $f_{\mathcal{Q}\times\mathcal{S}}(q,s) = f_{\mathcal{Q}}(q)\cdot f_{\mathcal{S}}(s;q)$ to denote the joint probability density of $q$ and $s$, and $f_{\mathcal{S}}(s) = \int_{q=0}^{1} f_{\mathcal{Q}}(q) f_{\mathcal{S}}(s;q)\,\mathrm{d}q$ to denote the marginal probability density of $s$.

We argue that $\mathsf{E}\Big[2q - m - s \;\Big|\; s, q \in [s,m]\Big] \leqslant 0$. Intuitively, this is because the posterior distribution of $q \in [s,m]$ conditioned on a fixed value of $s$ and on $q \in [s,m]$ is single-peaked with peak at $s$ by Assumptions 1 and 3. Hence, $\mathsf{E}\Big[q \;\Big|\; s, q \in [s,m]\Big] \leqslant (s+m)/2$. Formally, this can be viewed as

$$\begin{aligned}\mathsf{E}\Big[2q - m - s \;\Big|\; s, q \in [s,m]\Big] &= \frac{\int_{q=s}^{m} f(q;s)(2q-m-s)\,\mathrm{d}q}{\int_{q=s}^{m} f(q;s)\,\mathrm{d}q} \\ &\leqslant \frac{\frac{1}{m-s}\cdot\left(\int_{q=s}^{m} f(q;s)\,\mathrm{d}q\right)\cdot\left(\int_{q=s}^{m}(2q-m-s)\,\mathrm{d}q\right)}{\int_{q=s}^{m} f(q;s)\,\mathrm{d}q} = 0,\end{aligned}$$

where $f(q;s) = \frac{f_{\mathcal{Q}}(q)\cdot f_{\mathcal{S}}(s;q)}{f_{\mathcal{S}}(s)}$ denotes the probability density of true quality being $q$ conditioned on the score being $s$; the second transition is the integral Chebyshev inequality (Lemma 6), which holds because $f(q;s)$ is a non-increasing function of $q$ whereas $2q-m-s$ is a non-decreasing function of $q$;[6,7] and the final transition holds because the second integral in the numerator is $0$.

3. $m \leqslant q \leqslant s < h$. In this case, $|q-s| - |q - \mathrm{ULG}_T(s)| = m + s - 2q$. Due to the same reasoning as in Case 2, we have that $\mathsf{E}\Big[m + s - 2q \;\Big|\; m \leqslant q \leqslant s < h\Big] \leqslant 0$.

4. $m \leqslant s \leqslant q < h$. In this case, $|q-s| - |q - \mathrm{ULG}_T(s)| = m - s$. Due to the same reasoning as in Case 1, we have that $\mathsf{E}\Big[m - s \;\Big|\; m \leqslant s \leqslant q < h\Big] \leqslant -\Delta/8$.

Let $p_1, p_2, p_3, p_4$ respectively denote the total probabilities of the above four cases across all values of $k \in \{0,1,\ldots,T-1\}$, conditioned on $(q,s) \in \mathcal{D}^{\mathrm{same}}$. Then, $p_1+p_2+p_3+p_4 = 1$. Because $f_{\mathcal{S}}(a;b)/f_{\mathcal{S}}(b;a) \leqslant \gamma$ for all $a, b \in [0,1]$, it follows that $p_1 \geqslant p_2/\gamma$ and $p_4 \geqslant p_3/\gamma$. Hence, $p_1 + p_4 \geqslant (p_2+p_3)/\gamma$. Using $p_1+p_2+p_3+p_4 = 1$, we get $p_1 + p_4 \geqslant 1/(\gamma+1)$.

Combining the analysis from the four cases above, we have

$$\mathsf{E}\Big[|q-s| - |q - \mathrm{ULG}_T(s)| \;\Big|\; (q,s) \in \mathcal{D}^{\mathrm{same}}\Big] \leqslant -(p_1+p_4)\cdot\frac{\Delta}{8} \leqslant -\frac{\Delta}{8(\gamma+1)}. \tag{3.9}$$

**Analyzing $\mathcal{D}^{\mathrm{opp}}$.** Note that

$$\mathcal{D}^{\mathrm{opp}} = \cup_{k\in\{0,1,\ldots,T-1\}} \left\{(q,s) : (\ell(k) \leqslant q \leqslant m(k) \leqslant s \leqslant h(k)) \vee (\ell(k) \leqslant s \leqslant m(k) \leqslant q \leqslant h(k))\right\}.$$

Fix an arbitrary $k \in \{0,1,\ldots,T-1\}$; as before, write $\ell$, $m$, and $h$ while omitting the fixed $k$ in the argument. Once again, we analyze the desired expression $|q-s| - |q - \mathrm{ULG}_T(s)|$ conditioned on each

[6] To see why $f(q;s) = \frac{f_{\mathcal{Q}}(q)\cdot f_{\mathcal{S}}(s;q)}{f_{\mathcal{S}}(s)}$ is non-increasing in $q$, note that the denominator does not depend on $q$ whereas the numerator is equal to $f_{\mathcal{S}}(s;q)$ (Assumption 1), which is non-increasing in $q$ (Assumption 3).

[7] Technically, integral Chebyshev inequality requires non-negative functions, and $2q-(m+s)$ can be negative when $q < (m+s)/2$. However, one can equivalently separate out the $-(m+s)$ term, apply the integral Chebyshev inequality to $2q$, and recombine with the $-(m+s)$ term to achieve the same conclusion.

of the two cases in the above expansion of $\mathcal{D}^{\text{opp}}$ for this fixed $k$ separately. We will derive bounds that will hold regardless of the value of $k$, and, therefore, also conditional on $(q, s) \in \mathcal{D}^{\text{opp}}$ (i.e., aggregated across all $k$). Note that we still have $\text{ULG}_T(q) = \text{ULG}_T(s) = m$.

1. $\ell \leqslant q \leqslant m \leqslant s \leqslant h$: In this case, we have $|q - s| - |q - \text{ULG}_T(s)| = s - m$. Note that

$$\mathsf{E}\Big[s - m \;\Big|\; \ell \leqslant q \leqslant m \leqslant s \leqslant h\Big] \leqslant \Delta/4. \tag{3.10}$$

   This is because $s \in [m, m + \Delta/2]$ and, due to single-peakedness of the score model and $q \leqslant m$, it is at most $m + \Delta/4$ in expectation.

2. $\ell \leqslant s \leqslant m \leqslant q \leqslant h$: In this case, we have $|q - s| - |q - \text{ULG}_T(s)| = m - s$, and the same reasoning as above shows that

$$\mathsf{E}\Big[m - s \;\Big|\; \ell \leqslant s \leqslant m \leqslant q \leqslant h\Big] \leqslant \Delta/4. \tag{3.11}$$

Combining Equations (3.10) and (3.11) and aggregating over all $k \in \{0, 1, \ldots, T - 1\}$, we get that

$$\mathsf{E}\Big[|q - s| - |q - \text{ULG}_T(s)| \;\Big|\; (q, s) \in \mathcal{D}^{\text{opp}}\Big] \leqslant \Delta/4. \tag{3.12}$$

Finally, combining Equations (3.9) and (3.12), we have that

$$\begin{aligned}
&\mathsf{E}\Big[|q - s| - |q - \text{ULG}_T(s)| \;\Big|\; (q, s) \in \mathcal{D}\Big] \\
&\quad \leqslant \Pr\Big[(q, s) \in \mathcal{D}^{\text{same}} \;\Big|\; (q, s) \in \mathcal{D}\Big] \cdot \left(-\frac{\Delta}{8(\gamma + 1)}\right) + \Pr\Big[(q, s) \in \mathcal{D}^{\text{opp}} \;\Big|\; (q, s) \in \mathcal{D}\Big] \cdot \frac{\Delta}{4} \leqslant 0,
\end{aligned}$$

where the final transition holds because $\Pr\Big[(q, s) \in \mathcal{D}^{\text{same}}\Big] \geqslant 2(\gamma + 1) \cdot \Pr\Big[(q, s) \in \mathcal{D}^{\text{opp}}\Big]$ (Assumption 4), which is equivalent to

$$\Pr\Big[(q, s) \in \mathcal{D}^{\text{same}} \;\Big|\; (q, s) \in \mathcal{D}\Big] \geqslant 2(\gamma + 1) \cdot \Pr\Big[(q, s) \in \mathcal{D}^{\text{same}} \;\Big|\; (q, s) \in \mathcal{D}\Big].$$

This completes the proof. □

Ex-ante single-peakedness, ex-post single-peakedness, and probabilistic single-dippedness can be subsumed into a single property that captures a stronger form of symmetry, in which the noise in the score is symmetric and zero-mean.

**Definition 6** (Strongly Symmetric Score Model)**.** *We say that the score model $\mathcal{S}$ is strongly symmetric if $f_{\mathcal{S}}(s; q) = \ell(|s - q|)$ for some non-increasing function $\ell : \mathbb{R}_{\geqslant 0} \to \mathbb{R}_{\geqslant 0}$.*

Under a strongly symmetric score model, we have $\gamma = 1$ in Assumption 4 of Theorem 6, which means a constant of $2(\gamma + 1) = 4$ would be needed. However, using different techniques, we can show that even a constant of 3 suffices to obtain the same result under strong symmetry. This broadens the scope to include less concentrated score models.

**Theorem 7.** *Fix arbitrary $T \in \mathbb{N}$. Let $\mathcal{D}$, $\mathcal{D}^{\text{same}}$, and $\mathcal{D}^{\text{opp}}$ be defined as in Theorem 6. Assume the following regarding the true quality prior $\mathcal{Q}$ and the score model $\mathcal{S}$.*

1. *$\mathcal{Q}$ is uniform over $[0, 1]$;*
2. *$\mathcal{S}$ is strongly symmetric; and*
3. $\Pr\Big[(q, s) \in \mathcal{D}^{\text{same}}\Big] \geqslant 3 \cdot \Pr\Big[(q, s) \in \mathcal{D}^{\text{opp}}\Big]$.

*Then, the first implication of Corollary 4 holds. That is, uniform letter grading with $T$ grades is at least as good as numerical scoring when $\alpha_m > \alpha_d$, the converse holds when $\alpha_m < \alpha_d$, and the two have equal performance when $\alpha_m = \alpha_d$.*

*Proof.* As in the proof of Theorem 6, note that given Theorem 5, we only need to prove

$$\begin{aligned}
&\mathsf{E}\Big[|q - s| - |q - \mathrm{ULG}_T(s)| \;\Big|\; (q, s) \in \mathcal{D}\Big] \\
&\quad = \Pr[(q, s) \in \mathcal{D}^{\text{same}} \;\Big|\; (q, s) \in \mathcal{D}] \cdot \mathsf{E}\Big[|q - s| - |q - \mathrm{ULG}_T(s)| \;\Big|\; (q, s) \in \mathcal{D}^{\text{same}}\Big] \\
&\qquad + \Pr[(q, s) \in \mathcal{D}^{\text{opp}} \;\Big|\; (q, s) \in \mathcal{D}] \cdot \mathsf{E}\Big[|q - s| - |q - \mathrm{ULG}_T(s)| \;\Big|\; (q, s) \in \mathcal{D}^{\text{opp}}\Big] \\
&\quad \leqslant 0.
\end{aligned} \tag{3.13}$$

In the proof of Theorem 6, we analyzed the expected value of $|q - s| - |q - \mathrm{ULG}_T(s)|$ conditioned on both $(q, s) \in \mathcal{D}^{\text{same}}$ and $(q, s) \in \mathcal{D}^{\text{opp}}$ separately: the former was shown to be at most $-\frac{\Delta}{8(\gamma+1)}$ whereas the latter was shown to be at most $\frac{\Delta}{4}$, yielding the desired Equation (3.13) when $\Pr[(q, s) \in \mathcal{D}^{\text{same}}] \geqslant 2(\gamma + 1) \cdot \Pr[(q, s) \in \mathcal{D}^{\text{opp}}]$.

With strong symmetry (Assumption 2), we improve the former upper bound to $-\frac{\Delta}{12}$, which improves the sufficient condition to $\Pr[(q, s) \in \mathcal{D}^{\text{same}}] \geqslant 3 \cdot \Pr[(q, s) \in \mathcal{D}^{\text{opp}}]$. That is, our goal is to prove

$$\mathsf{E}\Big[|q - s| - |q - \mathrm{ULG}_T(s)| \;\Big|\; (q, s) \in \mathcal{D}^{\text{same}}\Big] \leqslant -\frac{\Delta}{12}.$$

Note that $\mathcal{D}^{\text{same}} = \cup_{k \in \{0,1,\ldots,T-1\}} \mathcal{D}^{\text{same}}_k$, where $\mathcal{D}^{\text{same}}_k = \mathcal{D}^{\text{same}} \cap [k\Delta, (k + 1)\Delta)^2$. We show that $\mathsf{E}\Big[|q - s| - |q - \mathrm{ULG}_T(s)| \;\Big|\; (q, s) \in \mathcal{D}^{\text{same}}_k\Big] \leqslant -\frac{\Delta}{12}$ for all $k \in \{0, 1, \ldots, T - 1\}$, which implies the desired result. Fix any $k \in \{0, 1, \ldots, T - 1\}$, and write $\ell = k\Delta$, $m = (k + 1/2)\Delta$, and $h = (k + 1)\Delta$.

Let us further partition $\mathcal{D}^{\text{same}}_k$ as $\mathcal{D}^{\text{same}}_{k,\text{low}} \cup \mathcal{D}^{\text{same}}_{k,\text{high}}$, where $\mathcal{D}^{\text{same}}_{k,\text{low}} = \{(q, s) : \ell \leqslant q, s < m\}$ (both the true quality and the score are lower than the midpoint) and $\mathcal{D}^{\text{same}}_{k,\text{high}} = \{(q, s) : m \leqslant q, s < h\}$ (both the true quality and the score are at least as high as the midpoint). Crucially, we note that

$$\mathsf{E}\Big[|q - s| - |q - m| \;\Big|\; (q, s) \in \mathcal{D}^{\text{same}}_{k,\text{low}}\Big] = \mathsf{E}\Big[|q - s| - |q - m| \;\Big|\; (q, s) \in \mathcal{D}^{\text{same}}_{k,\text{high}}\Big].$$

This is because the transformation $(q, s) \to (q', s')$, where $q' = m + (m - q)$ and $s' = m + (m - s)$, is a bijection mapping each point $(q, s) \in \mathcal{D}^{\text{same}}_{k,\text{low}}$ to a point $(q', s') \in \mathcal{D}^{\text{same}}_{k,\text{high}}$ with $|q - s| - |q - m| = |q' - s'| - |q' - m|$ as well as $f_{\mathcal{Q}\times\mathcal{S}}(q, s) = f_{\mathcal{Q}\times\mathcal{S}}(q', s')$; the last observation relies on $\mathcal{Q}$ being a uniform distribution (Assumption 1) and $\mathcal{S}$ being strongly symmetric (Assumption 2).

Hence, it is sufficient to show that

$$\mathsf{E}\Big[|q-s|-|q-m| \;\Big|\; (q,s) \in \mathcal{D}^{\text{same}}_{k,\text{low}}\Big] \leqslant -\frac{\Delta}{12}.$$

Next, we further partition $\mathcal{D}^{\text{same}}_{k,\text{low}}$ as $\mathcal{D}^{\text{same}}_{k,\text{low,inc}} \cup \mathcal{D}^{\text{same}}_{k,\text{low,dec}}$, where $\mathcal{D}^{\text{same}}_{k,\text{low,inc}} = \{(q,s) : \ell \leqslant q \leqslant s < m\}$ (the score is at least as much as the true quality) and $\mathcal{D}^{\text{same}}_{k,\text{low,dec}} = \{(q,s) : \ell \leqslant s \leqslant q < m\}$ (the score is at most as much as the true quality). Note that

$$\begin{aligned}
&\mathsf{E}\Big[|q-s|-|q-m| \;\Big|\; (q,s) \in \mathcal{D}^{\text{same}}_{k,\text{low}}\Big] \\
&\quad= \Pr\Big[(q,s) \in \mathcal{D}^{\text{same}}_{k,\text{low,inc}} \;\Big|\; (q,s) \in \mathcal{D}^{\text{same}}_{k,\text{low}}\Big] \cdot \mathsf{E}\Big[|q-s|-|q-m| \;\Big|\; (q,s) \in \mathcal{D}^{\text{same}}_{k,\text{low,inc}}\Big] \\
&\quad+ \Pr\Big[(q,s) \in \mathcal{D}^{\text{same}}_{k,\text{low,dec}} \;\Big|\; (q,s) \in \mathcal{D}^{\text{same}}_{k,\text{low}}\Big] \cdot \mathsf{E}\Big[|q-s|-|q-m| \;\Big|\; (q,s) \in \mathcal{D}^{\text{same}}_{k,\text{low,dec}}\Big].
\end{aligned}$$

First, we argue that

$$\Pr\Big[(q,s) \in \mathcal{D}^{\text{same}}_{k,\text{low,inc}} \;\Big|\; (q,s) \in \mathcal{D}^{\text{same}}_{k,\text{low}}\Big] = \Pr\Big[(q,s) \in \mathcal{D}^{\text{same}}_{k,\text{low,dec}} \;\Big|\; (q,s) \in \mathcal{D}^{\text{same}}_{k,\text{low}}\Big] = \frac{1}{2}.$$

This follows by noting the bijection from $\mathcal{D}^{\text{same}}_{k,\text{low,inc}}$ to $\mathcal{D}^{\text{same}}_{k,\text{low,dec}}$ given by $(q,s) \to (q',s')$, where $q' = m-(q-\ell)$ and $s' = m-(s-\ell)$; due to strong symmetry of $\mathcal{S}$ and $|q-s| = |q'-s'|$, we have $f_{\mathcal{Q}\times\mathcal{S}}(q,s) = f_{\mathcal{Q}\times\mathcal{S}}(q',s')$.

Next, recall that in the proof of Theorem 6 (Case 2 in the analysis of $\mathcal{D}^{\text{same}}$), we had already argued

$$\mathsf{E}\Big[|q-s|-|q-m| \;\Big|\; (q,s) \in \mathcal{D}^{\text{same}}_{k,\text{low,dec}}\Big] = \mathsf{E}[2q-s-m | \ell \leqslant s \leqslant q < m] \leqslant 0.$$

Hence, we have

$$\mathsf{E}\Big[|q-s|-|q-m| \;\Big|\; (q,s) \in \mathcal{D}^{\text{same}}_{k,\text{low}}\Big] \leqslant \frac{1}{2} \cdot \mathsf{E}\Big[|q-s|-|q-m| \;\Big|\; (q,s) \in \mathcal{D}^{\text{same}}_{k,\text{low,inc}}\Big],$$

which means it is sufficient to argue

$$\mathsf{E}\Big[|q-s|-|q-m| \;\Big|\; (q,s) \in \mathcal{D}^{\text{same}}_{k,\text{low,inc}}\Big] = \mathsf{E}\Big[s-m \;\Big|\; \ell \leqslant q \leqslant s < m\Big] \leqslant -\frac{\Delta}{6}.$$

Note that

$$\begin{aligned}
&\mathsf{E}\Big[s-m \;\Big|\; \ell \leqslant q \leqslant s < m\Big] \\
&\quad= -\frac{\int_{q=\ell}^{m}\int_{s=q}^{m}(m-s) f_{\mathcal{S}}(s;q)\,\mathrm{d}s\,\mathrm{d}q}{\Pr[\ell \leqslant q \leqslant s < m]} && (\mathcal{Q} \text{ is uniform}) \\
&\quad\leqslant -\frac{\int_{q=\ell}^{m}\frac{1}{m-q}\left(\int_{s=q}^{m}(m-s)\,\mathrm{d}s\right)\cdot\left(\int_{s=q}^{m} f_{\mathcal{S}}(s;q)\,\mathrm{d}s\right)\mathrm{d}q}{\Pr[\ell \leqslant q \leqslant s < m]} && (\text{Lemma 6}) \\
&\quad= -\frac{1}{2}\frac{\int_{q=\ell}^{m}(m-q)\Pr[s \in [q,m]]\,\mathrm{d}q}{\Pr[\ell \leqslant q \leqslant s < m]} \\
&\quad\leqslant -\frac{1}{2}\frac{\frac{2(m-\ell)}{3}\cdot\int_{q=\ell}^{m}\Pr[s \in [q,m]]\,\mathrm{d}q}{\Pr[\ell \leqslant q \leqslant s < m]} && (\text{Lemma 9}) \\
&\quad= -\frac{1}{2}\frac{\frac{2(m-\ell)}{3}\cdot\Pr[\ell \leqslant q \leqslant s < m]}{\Pr[\ell \leqslant q \leqslant s < m]} = -\frac{\Delta}{6},
\end{aligned}$$

as needed. Here, in the application of Lemma 9 in the fourth transition, we use the fact that $g(q) = \Pr[s \in [q, m]] = \int_{s=q}^{m} f_{\mathcal{S}}(s; q)\,\mathrm{d}s$ is a concave function and $g(m) = 0$. To see concavity, note that strong symmetry of $\mathcal{S}$ means that there is a distribution with probability density $z$ such that $f_{\mathcal{S}}(s; q) = z(s - q)$. Then, $g(q) = \int_{s=q}^{m} z(s - q)\,\mathrm{d}s = \int_{x=0}^{m-q} z(x)\,\mathrm{d}x$. Hence, $g'(q) = -z(m - q)$ and $g''(q) = z'(m - q)$. Due to the single-peakedness of $\mathcal{S}$, we have that $z'(x) \leqslant 0$ for all $x \geqslant 0$, so $g''(q) \leqslant 0$, which proves concavity of $g$. □

We remark that in the proof of Theorem 7, we only really need strong symmetry for pairs of true qualities and scores that belong to the same letter grade interval.

## 3.4 Experiments

In the previous section, we proved that when $\mathcal{Q}$ is uniformly distributed and the variance of the score model is small, we can conclude that the first implication of Corollary 4 holds. In this section, we empirically compare numerical scoring and uniform letter grading while relaxing these assumptions.

First, it is widely believed that students' true qualities, at least in large classes, are normally distributed based on the evidence that

> *...exam scores tend to be normally distributed for well-constructed, norm-referenced, multiple choice tests...[Wedell et al., 1989].*

Hence, we empirically study the case where $\mathcal{Q}$ is normally distributed, truncated to $[0, 1]$. We also consider cases where the score is not necessarily concentrated around the true quality. Finally, our analysis was limited to two evaluations; in our experiments, we also consider more than two evaluations. When a student participates in $r$ sequential evaluations, after each evaluation the student compares her "current" true quality to the obtained grade, and experiences (de)motivation that affects her effective true quality in the next evaluation. Formally, for $j \in [r]$, let $q_j$ and $s_j$ denote her effective true quality and score in evaluation $j$, respectively. Then, $s_j \sim \mathcal{S}(q_j)$ for each $j \in [r]$, and for $j \in [r - 1]$, we have:

$$q_{j+1} = \begin{cases} q_j + \alpha_m \cdot (B(s_j) - q_j), & \text{if } B(s_j) \geqslant q_j, \\ q_j - \alpha_d \cdot (q_j - B(s_j)), & \text{if } B(s_j) < q_j. \end{cases}$$

We measure the performance of a grading scheme by comparing the final true quality, $q_r$, to the initial true quality $q_1$, which extends the performance measure introduced in preliminaries for two evaluations:

$$\mathrm{perf}_F(B) \triangleq \mathsf{E}_{q\sim\mathcal{Q}, s\sim\mathcal{S}(q)}[q_r - q_1].$$

**Data generation.** For all the simulations, we compare numerical scoring (NS) to uniform letter grading ($\mathrm{ULG}_T$) with $T \in \{4, 8, 12, 16, 20\}$ grades. We scale the interval of grades to $[0, 100]$ to resemble percentage grades. We simulate $n = 5000$ students (average results are plotted with $95\%$ confidence intervals), where the initial true quality $q_1$ of each student is drawn i.i.d. from a truncated normal distribution capped to $[0, 100]$, with the underlying normal distribution characterized by mean $\mu$ and standard deviation $\sigma$. Given a true quality $q$ in an evaluation, the score $s$ is drawn from another truncated normal distribution capped to $[0, 100]$, with the underlying normal distribution characterized by mean $q$ and standard deviation $\gamma$.

**Results.** Figure 3.2 shows how the final quality improves (or degrades) with respect to the motivation coefficient (top) and the number of evaluations (bottom). In Figure 3.2a, the motivation coefficient takes values in $\{0, 0.1 \ldots, 0.9, 1\}$, the demotivation coefficient is set to $0.5$ and the number of evaluations is set to $r = 2$. We see that when $\alpha_m < \alpha_d$, numerical scoring is better than any uniform letter grading (and uniform letter grading with more grades is better than uniform letter grading with fewer grades), whereas when $\alpha_m > \alpha_d$, the opposite is true. Hence, it seems that the first implication of Corollary 4 continues to hold, even when the true quality is drawn from more realistic distributions. The comparison between uniform letter grading schemes with different numbers of grades is intuitive: uniform letter grading essentially converges to numerical scoring when $T$ goes to infinity, so larger $T$ should resemble numerical scoring more. The experiments show that this holds even with small values of $T$.

Going beyond our theoretical analysis for $r = 2$ evaluations, we consider the case where students participate in more than two evaluations. Surprisingly, as seen in Figures 3.2c and 3.2d, the comparison between numerical scoring and uniform letter grading flips completely with large values of $r$: numerical scoring becomes *worse* than uniform letter grading (and $\mathrm{ULG}_T$ becomes worse than $\mathrm{ULG}_{T'}$ for $T > T'$) when $\alpha_m < \alpha_d$, but *better* when $\alpha_d < \alpha_m$. This shows that the choice of the grading scheme depends not only on the comparison between the strengths of motivational and demotivational effects ($\alpha_m$ vs $\alpha_d$) but also, crucially, on the number of evaluations $r$. When $\alpha_m < \alpha_d$, with fewer evaluations (e.g., courses with fewer tests or curricula with fewer semesters), use of numerical scoring may be recommended, whereas with many evaluations (e.g., courses with frequent tests or curricula with many semesters), use of uniform letter grading with fewer letters may be more appropriate.

The transition between the regimes of few evaluations and many evaluations is even more surprising. As seen in Figure 3.2b, with $r = 4$ evaluations, numerical scoring seems to outperform uniform letter grading schemes regardless of the comparison between $\alpha_m$ and $\alpha_d$. Hence, in general, it is always best to simulate different grading schemes under the model and the number of evaluations of interest in order to pick a suitable grading scheme.

Finally, we observe that under numerical scoring, as the number of evaluations increases, the average student quality declines linearly when $\alpha_m < \alpha_d$ (Figure 3.2c) and improves linearly when $\alpha_m > \alpha_d$ (Figure 3.2d). This is expected because it can be shown that under numerical scoring, every evaluation changes the expected student quality by the same amount, which is proportional to $\alpha_m - \alpha_d$, leading to a linear decline or growth. In contrast, under uniform letter grading schemes with very few grades (small $T$), the average student quality seems to converge and remain stable as the number of evaluations increases, regardless of the comparison between $\alpha_m$ and $\alpha_d$. This can be explained due to the following stabilizing effect. Let $[\ell, h]$ be a letter grade interval and $m$ be its midpoint. Consider a student who starts with a true quality $q \in [\ell, h]$. The student is likely to receive a score $s$ in the same interval $[\ell, h]$ (so that $(q, s) \in \mathcal{D}$), and thus, a grade of $m$. This causes the true quality to update in a manner so that it gets closer to $m$ after which the student experiences very little motivation or demotivation due to receiving a grade that is almost equal to her true quality. Of course, the effect is more pronounced when $T$ is small, so letter grade intervals are large compared to the variance of the score model.

Moreover, we present additional experimental results, which show the impact of varying the mean $\mu$ of the true quality distribution (Figure 3.3), the standard deviation $\sigma$ of the true quality distribution (Figure 3.4), and the standard deviation $\gamma$ of the score distribution (Figure 3.5).[8]

[8]Technically, these are the mean and the standard deviations of the respective underlying normal distributions before truncation.

Overall, the mean true quality $\mu$ has little impact on the performance of different grading schemes. Similarly, the standard deviation $\sigma$ of the true quality prior also does not significantly affect the performance of the grading schemes, but somewhat strikingly, it has a dramatic impact on the performance of $\mathrm{ULG}_4$ (uniform letter grading with $4$ grades).

The impact of the standard deviation $\gamma$ of the score distribution is more significant, since as $\gamma$ increases, the performance of the different grading schemes becomes more similar. However, we see that even for quite large values of of $\gamma$, our theoretical results seem to hold. In particularly, we see that when two evaluations are taken place, numerical scoring is better when $\alpha_m < \alpha_d$ whereas uniform letter grading is better when $\alpha_m > \alpha_d$. This observation is quite encouraging since it probably indicates that our results can be extended for cases where the score is not very well-concentrated around the true quality.

## 3.5 Discussion

Our work takes the first step towards proposing a statistical model of the psychological impact of letter grading schemes on student performance in sequential evaluations and using it to compare uniform letter grading schemes to numerical scoring. Obviously, we model one specific (de)motivational effect and future work must combine it with many other well-known effects in educational settings. Nevertheless, we view our work as a stepping stone and outline several appealing extensions below.

**Beyond midpoint grading.** In our model, we assume that if all the scores from an interval $[\ell, u]$ are mapped to the same grade, they are effectively mapped to the midpoint grade $(\ell+u)/2$. This is a common method in practice of converting letter grades to percentages [University of Western Ontario, 2022, Victoria University of Wellington, 2022], but other values within the range $[\ell, u]$ are also sometimes used [University of Waterloo, 2022].

**Non-uniform letter grading.** Our analysis is limited to *uniform* letter grading schemes, which are used less often in practice. It would be interesting to extend our analysis to non-uniform letter grading schemes. More broadly, in our model, the grading scheme maps the score to a grade from $[0, 1]$, which allows a student to compare the grade to their true quality, which is also from $[0, 1]$. If the grade is instead in a different numerical range, the model can be easily extended by renormalization (for example, grade point averages lie in $[0, 4]$, which is often mapped to $[0, 1]$ by dividing by $4$). How can our model be extended to incorporate truly non-numeric grades (e.g., A, B, etc.) without converting them to numeric grades somehow (e.g., 4, 3.7, etc.)?

**Non-linear (de)motivation.** Our model assumes that the increase or decrease in the true quality is linear in the difference between the received grade and true quality. Evidence from prospect theory suggests that motivational effects from positive outcomes are typically concave (diminishing rewards) while demotivational effects from negative outcomes are typically convex (increasing losses) [Kahneman and Tversky, 1979]. It would be interesting to extend our theoretical results to nonlinear (de)motivational effects.

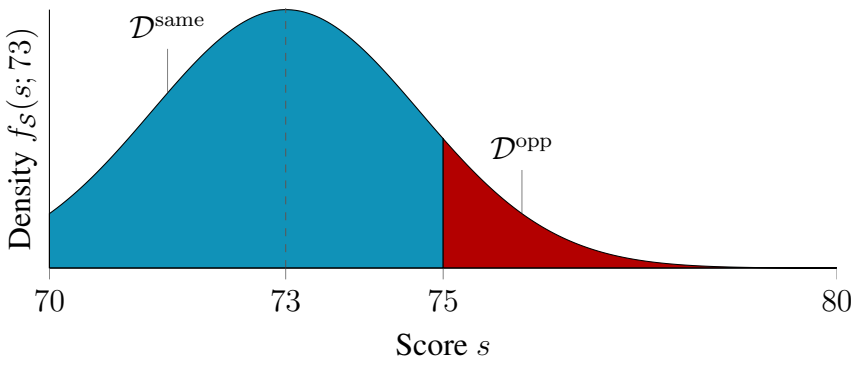


(a) With a small value of $\gamma$, one can see that within the grade interval $[70, 80]$ containing the true quality $q = 73$, the probability of the score being on the same side of the midpoint as the true quality (i.e., in $[70, 75]$) is significantly higher than the probability of it being on the opposite side of the midpoint (i.e., in $[75, 80]$). The former region contributes to $\mathcal{D}^{\text{same}}$ while the latter contributes to $\mathcal{D}^{\text{opp}}$. Their difference is the most pronounced when the true quality is near the interval endpoints (e.g., $q \approx 70, 80$) and gradually vanishes when it is near the midpoint (e.g., $q \approx 75$). In expectation over the true quality, one can still expect $\Pr[(q, s) \in \mathcal{D}^{\text{same}}]$ to be sufficiently higher than $\Pr[(q, s) \in \mathcal{D}^{\text{opp}}]$, satisfying the conditions in Theorem 6 and Theorem 7.

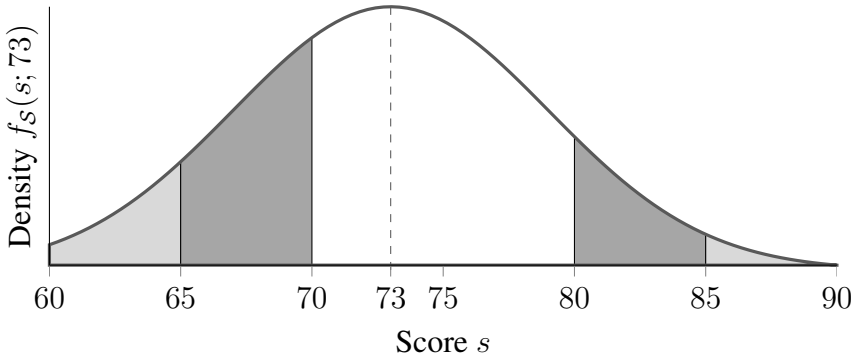


(b) Due to single-peakedness of the score distribution, the expected score in any interval lower than the interval containing the true quality $q = 73$ is at least its midpoint (e.g., the expected score subject to the score being in $[60, 70]$ is at least 65). In contrast, the expected score in any interval higher than the interval containing the true quality $q = 73$ is at most its midpoint (e.g., the expected score subject to the score being in $[80, 90]$ is at most 85). This observation is used at the end of the proof of Theorem 5.

Figure 3.1: Both figures show the probability density function of the score distribution $\mathcal{S}(q)$ when the true quality is $q = 73$. The distribution is a truncated normal distribution with mean $q = 73$, and standard deviation $\gamma = 1.7$ (top figure) and $\gamma = 6$ (bottom figure). The top figure conveys the intuition behind the conditions in Theorem 6 and Theorem 7, which assume $\Pr[(q, s) \in \mathcal{D}^{\text{same}}]$ to be sufficiently higher than $\Pr[(q, s) \in \mathcal{D}^{\text{opp}}]$. The bottom figure conveys the intuition behind the observation used at the end of the proof of Theorem 5.

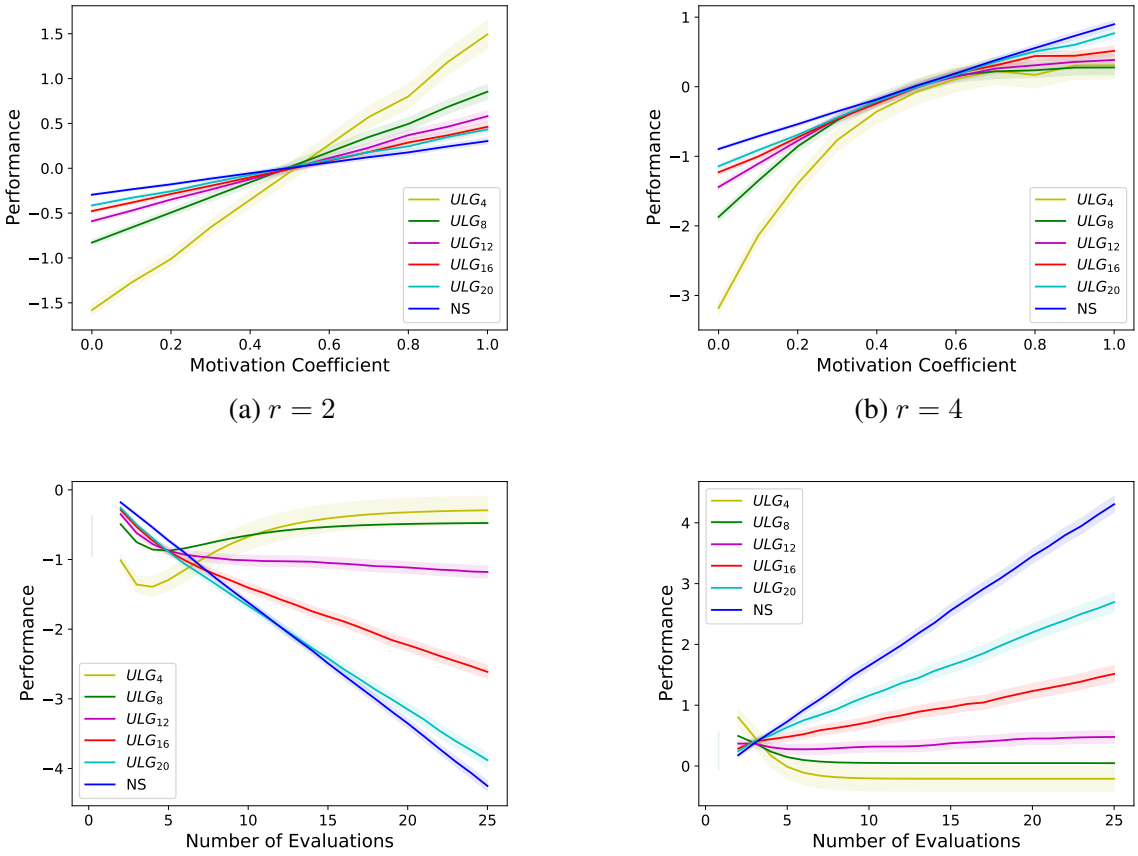


(a) $r = 2$

(b) $r = 4$

(c) $\alpha_m = 0.2$

(d) $\alpha_m = 0.8$

Figure 3.2: Performance of numerical scoring and different uniform letter grading schemes, with $\mu = 65$, $\sigma = 12$, $\gamma = 1.5$ and $\alpha_d = 0.5$ over different motivation coefficients (top) and number of evaluations (bottom). $95\%$ confidence intervals are shown.

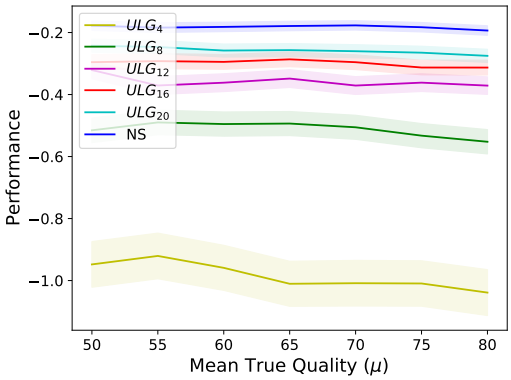


(a) $\alpha_m = 0.2, r = 2$

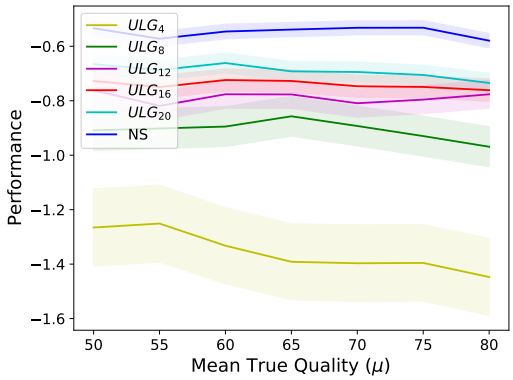


(b) $\alpha_m = 0.2, r = 4$

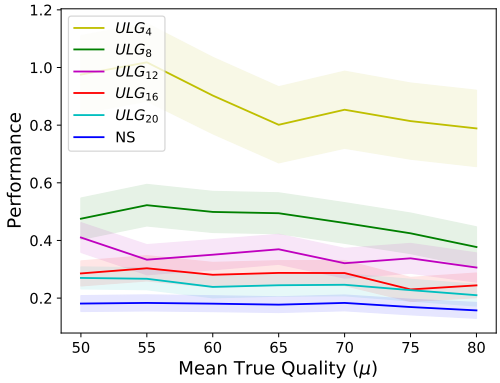


(c) $\alpha_m = 0.8, r = 2$

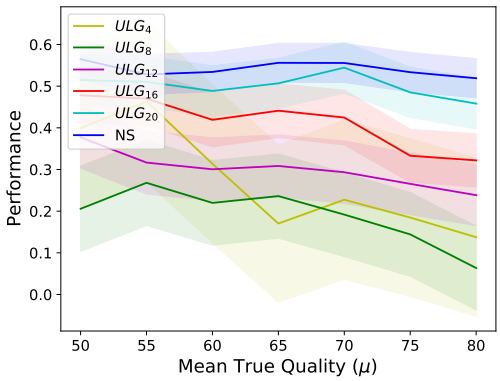


(d) $\alpha_m = 0.8, r = 4$

Figure 3.3: Performance of numerical scoring and different uniform letter grading schemes, with $\sigma = 12$, $\gamma = 1.5$ and $\alpha_d = 0.5$, over different values of $\mu$.

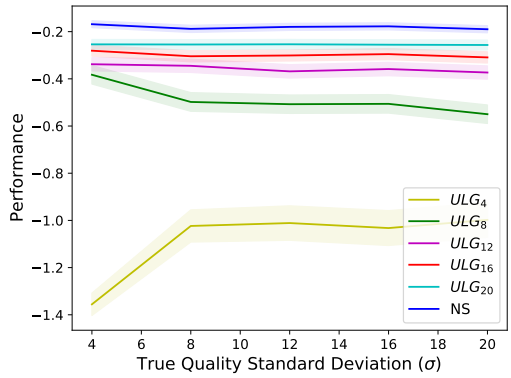


(a) $\alpha_m = 0.2, r = 2$

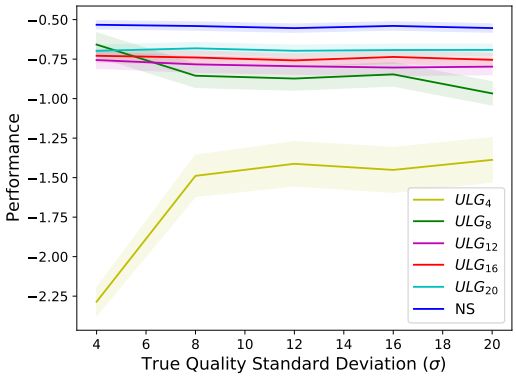


(b) $\alpha_m = 0.2, r = 4$

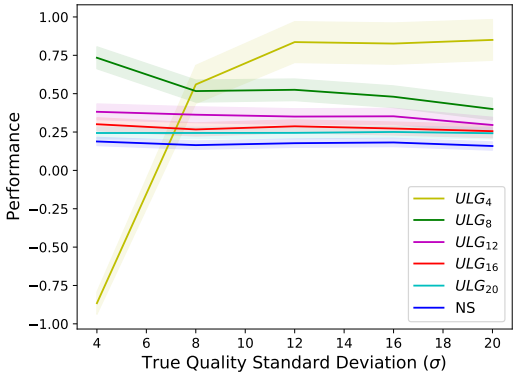


(c) $\alpha_m = 0.8, r = 2$

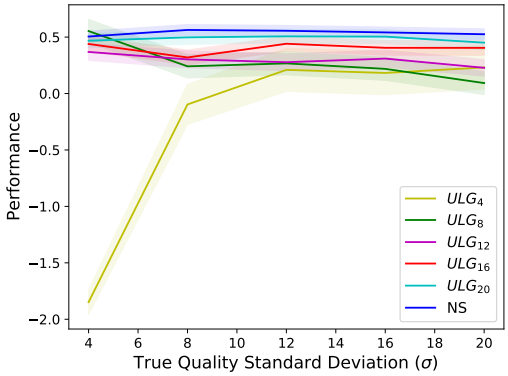


(d) $\alpha_m = 0.8, r = 4$

Figure 3.4: Performance of numerical scoring and different uniform letter grading schemes, with $\mu = 65$, $\gamma = 1.5$ and $\alpha_d = 0.5$, over different values of $\sigma$.

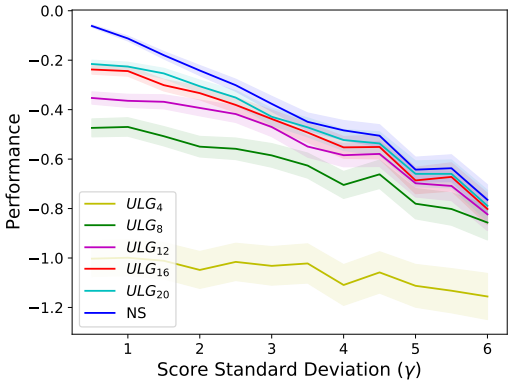


(a) $\alpha_m = 0.2, r = 2$

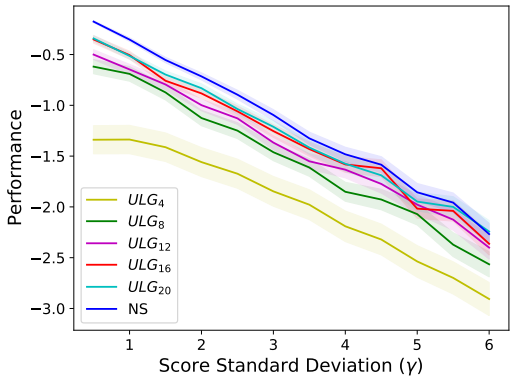


(b) $\alpha_m = 0.2, r = 4$

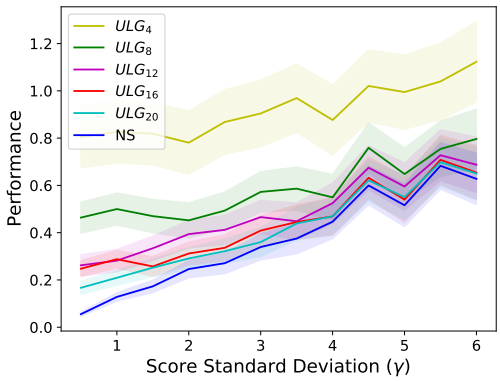


(c) $\alpha_m = 0.8, r = 2$

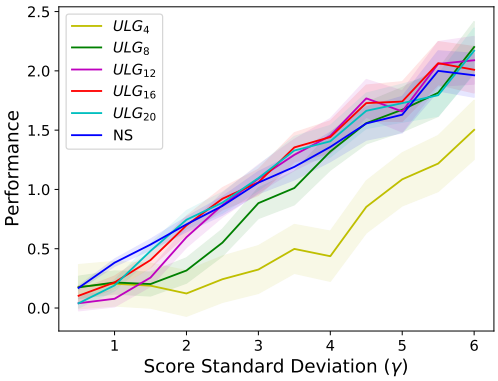


(d) $\alpha_m = 0.8, r = 4$

Figure 3.5: Performance of numerical scoring and different uniform letter grading schemes, with $\mu = 65$, $\sigma = 12$ and $\alpha_d = 0.5$, over different values of $\gamma$.

# Part II

# Civic Participation

# Chapter 4

# Fairness and Representation in Sortition

## 4.1 Introduction

Most people think of democracy as synonymous with elections. But that has not always been the case: from the inception of democracy in ancient Athens until the American and French revolutions, democracy had typically been associated with random selection of representatives Van Reybrouck [2016], a paradigm known as *sortition*.

These days, sortition is mainly practiced in the form of *citizens' assemblies* — randomly selected groups of people who deliberate on central questions, with the goal of generating recommendations and informing policy. The impact and prevalence of citizens' assemblies around the world have motivated computational work on how to fairly and transparently select assembly members [Flanigan et al., 2021a, 2020, 2021b]. But there are signs that sortition is becoming even more widely accepted, including its recent institutionalization in Belgium, where permanent sortition-based bodies are now working alongside the parliaments of the German-speaking region and the Brussels region. In light of this progress, it may only be a matter of time until one of the many blueprints for sortition-based democracy [Gastil and Wright, 2019] is implemented at the level of an entire country.

The excitement about sortition is driven by several appealing qualities, which are seen as providing solutions to some of the problems plaguing electoral democracy. We briefly present two of them in the context of *uniform selection*, which selects a uniformly random panel and is considered to be the ideal sortition method [Engelstad, 1989].

- *Descriptive representation*: A panel selected uniformly at random is likely to reflect the composition of the population from which it was drawn. Representation lends *legitimacy* to the process [Parker, 2011, Fishkin, 2018], as individuals are able to identify some panelists who are similar to themselves.[1]

- *Fairness*: Under uniform selection, each individual has an equal chance to participate. Political theorists have argued that this quality realizes philosophical ideals like equality of opportunity and allocative justice [Stone, 2011].

[1] Another motive for representation is epistemic: Landemore [2013] and others argue that a diversity of opinions leads to better decisions.

By any reasonable measure of the fairness of selection probabilities — e.g., the minimum selection probability of any individual [Flanigan et al., 2021a] — uniform selection achieves perfect fairness, as selection probabilities are equalized. We ask: *Is uniform selection also representative in a rigorous sense?* If we had an analogous measure of representation, we would be able to explore whether this is the case. But quantifying representation poses a conceptual challenge.

In this chapter, we aim to address the following questions:

1. *What is a formal definition of representation of a population?*
2. *To what extent does uniform selection satisfy representation?*
3. *Is it possible to design selection procedures that enhance representation guarantees while maintaining fairness?*

## 4.2 Model

For all $t \in \mathbb{N}$, define $[t] = \{1, \ldots, t\}$. Let $N = [n]$ be the set that indexes the underlying population. A *panel* $P$ is a subset of the population. Let $\mathcal{S}_k(N)$ denote the set of all subsets of $N$ of size $k$. (We omit $N$ when it is clear form the context.) The population lies in a metric space with distance $d$, which we view as the *representation metric* discussed earlier. For each $i, j \in N$, $d(i, j)$ denotes the distance between $i$ and $j$; $d$ is a metric if the following properties are satisfied: (a) $d(i, j) \geqslant 0$, and $d(i, j) = 0$ if and only if $i = j$, (b) $d(i, j) = d(j, i)$, and (c) for each $i, j, \ell \in N$, $d(i, \ell) + d(\ell, j) \geqslant d(i, j)$. The last property is known as the triangle inequality. An instance of our problem is given by the underlying population along with distances as defined by $d$; hereinafter, we simply denote such an instance by $d$.

Given a panel $P$ of size $k$ and a positive integer $q \in [k]$, the *$q$-cost* of individual $i$ for $P$, denoted by $c_q(i, P; d)$, is equal to the distance of $i$ from her $q$-th closest representative in $P$. Note that for $q = 1$, we have $c_1(i, P; d) = \min_{j \in P} d(i, j)$ and for $q = k$ we get $c_k(i, P; d) = \max_{j \in P} d(i, j)$. Let $\mathsf{top}_q(i, P; d)$ be the set of $q$ closest members of $P$ to $i$ (with ties broken arbitrarily). The *$q$-social cost* of panel $P$ is given by $\mathrm{SC}_q(P; d) = \sum_{i \in N} c_q(i, P; d)$, i.e., the sum of the $q$-costs over all individuals. (Observe that if $q = 1$ then $\mathrm{SC}_q(P; d)$ is the standard $k$-medians clustering objective evaluated for centers $P$.) We omit $d$ from the notation when it is clear from the context.

In this setting, a *selection algorithm* $\mathcal{A}_{k,q}$ parameterized by $k$ and $q$ takes as input the metric $d$ and outputs a distribution over all panels of size $k$. We are especially interested in the *uniform selection* algorithm, denoted by $\mathcal{U}_k$, that always outputs a uniform distribution over $\mathcal{S}_k$, regardless of $q$. In other words, it does not take into account the underlying metric space or $q$, but instead outputs a committee of size $k$ chosen uniformly at random. We now formally introduce the fairness a selection algorithm.

**Fairness:** One appealing property of uniform selection is that each individual is selected to be part of the panel with probability exactly equal to $k/n$, i.e., $\Pr[i \in \mathcal{U}_k] = \frac{k}{n}$. In particular, all individuals have an equal chance of being chosen. We call this property *perfect fairness*; in general an algorithm $\mathcal{A}_{k,q}$ provides perfect fairness when for each instance $d$, it ensures that $\min_{i \in N} \Pr[i \in \mathcal{A}_{k,q}] = k/n$.

When perfect fairness is too restrictive, we relax this constraint by allowing individuals to be selected with probability less than $k/n$. In this case, the fairness of an algorithm $\mathcal{A}_{k,q}$ is the worst-case ratio

of the minimum probability of an individual to be selected by the algorithm and the ideal selection probability of $k/n$. Formally,

$$\mathsf{fairness}_q(\mathcal{A}_{k,q}) = \inf_d \frac{\min_{i \in N} \Pr[i \in \mathcal{A}_{k,q}(d)]}{k/n}.$$

In the next two sections, we take two different approaches for formally defining the representation of a selection algorithm.

## 4.3 Representation Through $q$-Social Cost

In this section, we measure the representation of a selection algorithm by taking into account the expected $q$-social cost. In a high level, we consider a panel to be a good representative of the whole population when its $q$-social cost is not much larger than the best possible. In other words, a selection algorithm $\mathcal{A}_{k,q}$ that outputs a distribution over the different committees of size $k$ provides good representation when the expected $q$-social cost of the panel is similarly small. More formally, we define the *representation* of a selection algorithm $\mathcal{A}_{k,q}$ as the worst-case ratio of the minimum possible social $q$-cost of any panel and the expected $q$-social cost of the panel chosen by $\mathcal{A}_{k,q}$ over all possible instances; i.e.,

$$\mathsf{repr}_q(\mathcal{A}_{k,q}) = \inf_d \frac{\min_{P' \in \mathcal{S}_k(N)} \mathrm{SC}_q(P'; d)}{\mathbb{E}[\mathrm{SC}_q(\mathcal{A}_{k,q}(d)]}.$$

### 4.3.1 Representation with Perfect Fairness for $q > k/2$

We begin by considering the case that $q > k/2$. We show that in this case uniform selection is asymptotically optimal with respect to representation among all selection algorithms that are perfectly fair. Moreover, the representation of uniform selection is constant for any $q = c \cdot k$ for $1/2 < c < 1$.

**Theorem 8.** *For $q > k/2$, uniform selection satisfies $\mathsf{repr}_q(\mathcal{U}_k) \geqslant \frac{1}{2} \cdot \frac{k-q+1}{k}$.*

A crucial property that we exploit in this section is that the $q$-costs of the individuals satisfy the triangle inequality when $q > k/2$. This observation was first made by [Caragiannis et al., 2022]; we present the lemma below for completeness.

**Lemma 10.** *For $q > k/2$, individuals $i, j \in N$, and a panel $P$, $c_q(i, P; d) + c_q(j, P; d) \geqslant d(i, j)$.*

*Proof.* Let $T_i = \mathsf{top}_q(i, P)$ and $T_j = \mathsf{top}_q(i, P)$ be the $q$ closest neighbors of $i$ and $j$, respectively, in the panel $P$. As $|T_i| = |T_j| > k/2$, there exists an individual $k \in T_i \cap T_j$. Therefore,

$$d(i, j) \leqslant d(i, k) + d(k, j) \leqslant c_q(i, P) + c_q(j, P). \qquad \square$$

We use this observation to lower bound the social cost of the optimal committee.

**Lemma 11.** *For $q > k/2$, the $q$-social cost of the optimal panel $P^*$ is at least*

$$\mathsf{SC}_q(P^*; d) \geqslant \frac{1}{2(n-1)} \cdot \sum_{i \in N} \sum_{j \in N \setminus \{i\}} d(i, j).$$

*Proof.* By applying Lemma 10 for all pairs of individuals $(i, j)$, we get

$$\sum_{i \in N} \sum_{j \in N \setminus \{i\}} \Big( c_q(i, P^*) + c_q(j, P^*) \Big) \geqslant \sum_{i \in N} \sum_{j \in N \setminus \{i\}} d(i, j).$$

The $q$-cost of each person appears exactly $2(n-1)$ times on the left hand side. Thus,

$$\mathsf{SC}_q(P^*; d) = \sum_{i \in N} c_q(i, P^*) \geqslant \frac{1}{2(n-1)} \sum_{i \in N} \sum_{j \in N \setminus \{i\}} d(i, j). \qquad \square$$

We are now ready to prove the theorem.

*Proof of Theorem 8.* For any committee $P$ of size $k$,

$$\begin{aligned} c_q(i, P) = \min_{q' \in \{q, \ldots, k\}} c_{q'}(i, P) &\leqslant \frac{1}{k-q+1} \sum_{q' \in [q,k]} c_{q'}(i, P) \\ &\leqslant \frac{1}{k-q+1} \sum_{q' \in [1,k]} c_{q'}(i, P) = \frac{1}{k-q+1} \sum_{j \in P} d(i, j), \end{aligned}$$

where in the first inequality the minimum is upper bounded by the average. Therefore, the expected social cost of uniform selection is at most

$$\begin{aligned} \mathsf{E}[\mathsf{SC}_q(\mathcal{U}_k(N))] &= \sum_{i \in N} \mathsf{E}_{P \sim \mathcal{U}_k} [c_q(i, P)] \\ &\leqslant \frac{1}{k-q+1} \sum_{i \in N} \mathsf{E}_{P \sim \mathcal{U}_k} \Big[ \sum_{j \in P} d(i, j) \Big] \\ &= \frac{1}{k-q+1} \sum_{i \in N} \sum_{j \in N \setminus \{i\}} d(i, j) \cdot \Pr_{P \sim \mathcal{U}_k}[j \in P] \\ &= \frac{1}{k-q+1} \sum_{i \in N} \sum_{j \in N \setminus \{i\}} d(i, j) \cdot \frac{k}{n}. \end{aligned}$$

By Lemma 11 and the upper bound shown above, we have

$$\mathsf{repr}_q(\mathcal{U}_k) \geqslant \frac{\frac{1}{2(n-1)} \sum_{i \in N} \sum_{j \in N \setminus \{i\}} d(i, j)}{\frac{1}{k-q+1} \cdot \frac{k}{n} \cdot \sum_{i \in N} \sum_{j \in N \setminus \{i\}} d(i, j)} = \frac{1}{2} \cdot \frac{n}{n-1} \cdot \frac{k-q+1}{k} \geqslant \frac{1}{2} \cdot \frac{k-q+1}{k}. \qquad \square$$

In the proof above, the only property of uniform selection we use is that the marginal probabilities are equal to $k/n$. Hence, this lower bound also holds for any perfectly fair selection algorithm.

We next establish an upper bound on the representation of any perfectly fair selection algorithm. It shows that the lower bound of Theorem 8 is tight up to a factor of 4.

**Theorem 9.** *For any $q > k/2$, every selection algorithm $\mathcal{A}_{k,q}$ with $\mathsf{fairness}(\mathcal{A}_{k,q}) = 1$ satisfies $\mathsf{repr}_q(\mathcal{A}_{k,q}) \leqslant 2 \cdot \frac{k-q+1}{k+1}$.*

*Proof.* First, note that if $q = \frac{k+1}{2}$, the statement trivially holds, since $\mathsf{repr}_q(\mathcal{A}_{k,q}) \leqslant 2 \cdot \frac{k-q+1}{k+1} = 1$ which is true for any algorithm. Thus, we assume $q > \frac{k+1}{2}$. Consider an instance with $n = k+1$ individuals where $k-q+1$ individuals are located at 0 and $q$ individuals are at 1, denoted by $N_0$ and $N_1$, respectively. Any committee of size $k$ leaves one person out of the committee, and for each individual $i \in N$, this happens with probability of

$$\Pr_{P\sim\mathcal{A}_{k,q}}[i \notin P] = 1 - \Pr_{P\sim\mathcal{A}_{k,q}}[i \in P] = (1 - \tfrac{k}{k+1}) = \tfrac{1}{k+1}.$$

Individuals in $N_0$ will always have a $q$-cost of 1, because $k+1-q < \frac{k+1}{2} < q$ individuals are located there. Therefore, $\mathsf{E}_{P\sim\mathcal{A}_{k,q}}[\sum_{i\in N_0} c_q(i,P)] = |N_0|$. For individuals in $N_1$, their $q$-cost is 1 if and only if less than $q$ individuals are selected from $N_1$, i.e., the single person left out of the committee is located at 1. This event happens with probability

$$\Pr_{P\sim\mathcal{A}_{k,q}}[\textstyle\bigcup_{i\in N_1}(i \notin P)] = \sum_{i\in N_1} \Pr_{P\sim\mathcal{A}_{k,q}}[i \notin P] = \tfrac{|N_1|}{k+1},$$

where the first equality comes from the fact that the events are disjoint (which holds because any committee leaves out exactly one individual). Therefore, $\mathsf{E}_{P\sim\mathcal{A}_{k,q}}[\sum_{i\in N_1} c_q(i,P)] = |N_1| \cdot \frac{|N_1|}{k+1}$, and the expected social cost of any perfectly fair algorithm is

$$\mathsf{E}_{P\sim\mathcal{A}_{k,q}}[\mathsf{SC}_q(P;d)] = |N_0| + |N_1| \cdot \frac{q}{k+1} \geqslant |N_0| + |N_1| \cdot \frac{1}{2}.$$

The optimal committee would leave out a person from $N_0$ and achieve a social cost of $|N_0|$. Therefore, the representation of any algorithm with perfect fairness is at most

$$\frac{|N_0|}{|N_0| + \frac{1}{2}\cdot|N_1|} \leqslant \frac{|N_0|}{\frac{1}{2}\cdot|N_0| + \frac{1}{2}\cdot|N_1|} = 2 \cdot \frac{k-q+1}{k+1}.$$ □

### 4.3.2 Representation with Relaxed Fairness for $q \leqslant k/2$

In stark contrast to the case of $q > k/2$, uniform selection and, more generally, any perfectly fair selection algorithm, cannot obtain bounded representation when $q \leqslant k/2$. In fact, the following theorem shows that selection algorithms with fairness strictly more than $\frac{q+(k \bmod q)}{k}$ (which itself is upper bounded by $(2q-1)/k$) suffer from unbounded representation.

**Theorem 10.** *For $q \leqslant k/2$, $\epsilon > 0$, and any selection algorithm $\mathcal{A}_{k,q}$ with $\mathsf{fairness}(\mathcal{A}_{k,q}) \geqslant \frac{q+(k \bmod q)}{k} + \epsilon$, $\mathsf{repr}_q(\mathcal{A}_{k,q})$ is 0.*

*Proof.* Assume that $n > 2 \cdot \max\{\sqrt{kq/\epsilon}, k+1\}$, and let $m = \lfloor k/q \rfloor - 1$. Consider the real line, and suppose there are sets of $q$ individuals at each position in $\{1, 2, \ldots, m\}$, denoted by $X_1, \ldots, X_m$, respectively, and the set of remaining $n - mq$ individuals, denoted by $X_{m+1}$, is at position $m+1$. The optimal panel $P^*$ would have at least $q$ people from each position, i.e., $|P^* \cap X_i| \geqslant q$ for all $i \in [m+1]$. The $q$-cost of each person for $P^*$ is 0 as at least $q$ people are selected from her own position. Hence, $\mathsf{SC}_q(P^*) = 0$.

Turning to the analysis of $\mathsf{A}_{k,q}$, we claim that

$$\mathsf{E}_{P\sim\mathsf{A}_{k,q}}[|X_{m+1} \cap P|] \geqslant q + (k \bmod q) + \epsilon. \tag{4.1}$$

To prove this, note that since each individual is included with a marginal probability of at least $\frac{q+(k \bmod q)}{n} + \epsilon$, we have

$$\begin{aligned}\mathsf{E}_{P\sim\mathsf{A}_{k,q}}[|X_{m+1} \cap P|] &\geqslant \left(\frac{q + (k \bmod q)}{n} + \epsilon\right)(n - mq) \\ &= q + (k \bmod q) - \frac{mq \cdot (q + (k \bmod q))}{n} + (n - mq) \cdot \epsilon\end{aligned}$$

We will show that the right hand side is at least $q + (k \bmod q) + \epsilon$. Because $mq < k$, $q + k \bmod q < 2q$, and $n - mq > n - k \geqslant n/2 + 1$, the right hand side is at least

$$q + (k \bmod q) - \frac{qk}{n/2} + n\epsilon/2 + \epsilon \geqslant q + (k \bmod q) + \epsilon,$$

where the inquality follows from our choice of $n > 2\sqrt{kq/\epsilon}$. This establishes Equation (4.1).

Now, as the panel size is $k$, it holds that $\sum_{i\in[m+1]} \mathsf{E}_{P\sim\mathsf{A}_{k,q}}[|X_i \cap P|] = k$. By Equation (4.1),

$$\sum_{i\in[m]} \mathsf{E}_{P\sim\mathsf{A}_{k,q}}[|X_i \cap P|] < k - (q + (k \bmod q)) - \epsilon = mq - \epsilon.$$

Therefore, there exists $i \in [m]$ such that $\mathsf{E}_{P\sim\mathsf{A}_{k,q}}[|X_i \cap P|] \leqslant q - \epsilon/q$. Using Markov's inequality,

$$\Pr_{P\sim\mathsf{A}_{k,q}}(|X_i \cap P| \geqslant q) \leqslant \tfrac{q-\epsilon/q}{q} \leqslant 1 - \epsilon.$$

Thus, with probability at least $\epsilon$, less than $q$ people are selected from position $i$, in which case the $q$-cost of each person in $X_i$ will be at least 1. Hence, $\mathsf{E}_{P\sim\mathsf{A}_{k,q}}[\mathsf{SC}_q(\mathsf{A}_{k,q})] \geqslant q\epsilon$ while $\mathsf{SC}_q(P^*) = 0$. □

Although bounded representation is not feasible with fairness slightly larger than $\frac{q}{k}$, we design a selection algorithm that, given an $\alpha$-representative panel, can achieve representation of at least $\frac{\alpha}{q+1}$ with fairness of $\frac{q}{k}$. In particular, starting from the optimal panel ($\alpha = 1$), it achieves $\frac{1}{q+1}$ representation. Although finding the optimal panel given the metric space is an NP-hard problem for most metric spaces, computing a constant-factor approximation (i.e., constant $\alpha$) is feasible in polynomial time [Kumar and Raichel, 2013].

Our Algorithm, RANDOMREPLACE$_q$, is given in Algorithm 2. It starts from a panel $P$, randomly selects a panel $S$ of $q$ individuals, and replaces individuals in $P$ with individuals in $S$ as follows. First, individuals in $S \cap P$ remain in the final panel. Then, for $i \in S \setminus P$, swap $i$ with one of its $q$-closest neighbors in the optimal panel $\mathsf{top}_q(i, P)$ that has not been replaced by the algorithm yet. The next theorem establishes the fairness and representation guarantees of RANDOMREPLACE.

**Theorem 11.** *For any $q \in [k]$ and any panel $P$ with $\mathsf{repr}_q(P) = \alpha$, we have that*

$$\mathsf{repr}_q(\text{RANDOMREPLACE}_q(P)) \geqslant \frac{\alpha}{q+1}$$

*and* $\mathsf{fairness}(\text{RANDOMREPLACE}_q) \geqslant \frac{q}{k}$.

**Algorithm 2** RANDOMREPLACE$_q$

**Input:** Panel $P$ with $\mathsf{repr}_q(P) = \alpha$
**Output:** Randomly selected panel by replacing at most $q$ individuals of $P$

1: Pick $S \in \mathcal{S}_q$ uniformly at random
2: Set $P_S \leftarrow P$ and $\overline{S} \leftarrow S \setminus P$
3: **for** $i \in \overline{S}$ **do**
4: Pick an arbitrary $j_i \in \mathsf{top}_q(i, P) \setminus S$
5: $P_S \leftarrow P_S \cup \{i\} \setminus \{j_i\}$
6: **end for**
7: **return** $P_S$

*Proof.* Let $S \subseteq N$ be a set of size $q$ chosen uniformly at random. We denote with $P_S$ the panel that is returned from the algorithm. First, we show that Line 4 of the algorithm is valid. The algorithm reaches this line when it considers $i \in S \setminus P_S$, meaning that $i$ is not included in the panel $P_S$ and therefore is not included in $P$. Hence, $\mathsf{top}_q(i, P) \setminus S$ cannot be empty since $|\mathsf{top}_q(i, P)| = q$, $|S| = q$ and there exists $i$ in $S$ but not in $\mathsf{top}_q(i, P)$.

We see that every individual in $S$ is included in $P_S$ as in Line 5 the algorithms ensures that each such individual is included and is never excluded afterwards. Hence, as each individual is chosen in $S$ with probability at least $q/n$, we can see that $\mathsf{fairness}(\text{RANDOMREPLACE}_q) \geqslant q/k$.

Now, we prove that for any $S \in \mathcal{S}_q$ and any individual $i' \in N$,

$$c_q(i', P_S) \leqslant c_q(i', P) + \max_{i \in S} c_q(i, P) \leqslant c_q(i', P) + \sum_{i \in S} c_q(i, P). \tag{4.2}$$

The second inequality holds as the maximum is at most the sum. Therefore, we focus on the first inequality. If $c_q(i', P_S) \leqslant c_q(i', P)$, then it trivially holds. Otherwise, $c_q(i', P_S) > c_q(i', P)$. In this case, we can show that there exists some $i \in S \setminus P$ such that $d(i', i) \geqslant c_q(i', P_S)$ and $j_i \in \mathsf{top}_q(i', P)$. First, note if for each $i \in S \setminus P^*$, $j_i$ does not belong in $\mathsf{top}_q(i', P)$, then it is not possible that $c_q(i', P_S) > c_q(i', P)$, since $\mathsf{top}_q(i', P) \subseteq P_S$. Next, suppose for contradiction that for every $i$ that was included in $P_S$ when $j_i \in \mathsf{top}_q(i', P)$ was excluded from it, it holds that $d(i', i) < c_q(i', P)$. Then, in $P_S$ there are $|\mathsf{top}_q(i', P) \setminus \cup_{i \in S \setminus P}(\{j_i\} \cap \mathsf{top}_q(i', P))|$ individuals that have distance at most $c_q(i', P) < c_q(i', P_S)$ from $i'$ and $|\{i \in S \setminus P : j_i \in \mathsf{top}_q(i', P)\}|$ individuals that have distance less than $c_q(i', P_S)$ from $i'$. Note that

$$|\cup_{i \in S \setminus P}(\{j_i\} \cap \mathsf{top}_q(i', P))| = |\{i \in S \setminus P : j_i \in \mathsf{top}_q(i', P)\}|,$$

and hence we get that there are at least $|\mathsf{top}_q(i', P)| = q$ individuals in $P_S$ with distance strictly less than $c_q(i', P_S)$ from $i'$, which is a contradiction.

From the above observation , we have that

$$c_q(i', P_S) \leqslant d(i', i) \leqslant d(i', j_i) + d(j_i, i) \leqslant c_q(i', P) + c_q(i, P) \leqslant c_q(i', P) + \max_{i \in S} c_q(i, P)$$

where the penultimate inequality holds because $j_i \in \mathsf{top}_q(i', P)$ and $j_i \in \mathsf{top}_q(i, P)$ from the definition of $j_i$. This proves (4.2).

Summing (4.2) over all $i' \in N$, we have $\mathsf{SC}_q(P_S) \leqslant \mathsf{SC}_q(P) + n \cdot \sum_{i \in S} c_q(i, P)$. Taking the expectation of this equation with respect to $S$, and using the fact that $\Pr[i \in S] = q/n$, we have

$$\mathsf{E}[\mathsf{SC}_q(P_S)] \leqslant \mathsf{SC}_q(P) + n \cdot \sum_{i \in N} \frac{q}{n} \cdot c_q(i, P) = (q+1) \cdot \mathsf{SC}_q(P),$$

that is, $\frac{1}{q+1} \cdot \mathsf{E}[\mathsf{SC}_q(P_S)] \leqslant \mathsf{SC}_q(P)$. Combined with the assumption that $P$ is $\alpha$-representative, i.e., $\alpha \cdot \mathsf{SC}_q(P) \leqslant \min_{P' \in S_k(N)} \mathsf{SC}_q(P')$, we get $\mathsf{repr}_q(\textsc{RandomReplace}_q(P)) \leqslant \frac{\alpha}{k+1}$. □

The next theorem shows that when $q = \Omega(k)$, $\textsc{RandomReplace}_q$ attains an asymptotically optimal representation-fairness tradeoff.

**Theorem 12.** *For any $q \leqslant k/2$ such that $k \bmod q = 0$, if a selection algorithm $\mathcal{A}_{k,q}$ has $\mathsf{fairness}(\mathcal{A}_{k,q}) \geqslant q/k$, then $\mathcal{A}_{k,q}$ satisfies*

$$\mathsf{repr}_q(\mathcal{A}_{k,q}) \leqslant k/q^2.$$

*Proof.* Let $m = k/q$. Consider an instance with $n > 2k$ individuals on the real line, where one individual $i_0$ is located at 0, $\lceil n/m \rceil - 1$ people are at 1, and at least $\lfloor n/m \rfloor$ people are located at each position $j \in \{2, \ldots, m\}$. This way, there are at least $n/m - 1 = (n/k) \cdot q - 1 \geqslant 2q - 1 \geqslant q$ individuals located at each $j \in [m]$.

Any optimal panel would include $q$ individuals from each position $j \in [m]$, which results in a $q$-cost of 1 for $i_0$ and a $q$-cost of 0 for the rest. Hence, $\mathsf{SC}_q(P^*) = 1$. However, any selection algorithm with fairness of at least $q/k$ selects $i_0$ with probability of at least $q/n$. When $i_0$ is selected, there must exist a group $j \in [m]$ from which the algorithm selects at most $q - 1$ people, incurring a $q$-cost of 1 for at least $n/m$ people (person $i_0$ and at least $n/m - 1$ people at position $j$). Hence,

$$\mathsf{E}[\mathsf{SC}_q(\mathcal{A}_{k,q})] \geqslant \frac{q}{n} \cdot \frac{nq}{k} = \frac{q^2}{k},$$

which completes the proof. □

From the above theorem, it follows that Algorithm 2 achieves the highest possible fairness of $q/k$ subject to positive representation.

## 4.4 Representation Through the $q$-Core

In the previous section, we use the $q$-social cost to measure how much a panel represents a given population. In this section, we take a step back and ask what it really means a panel to reflect the composition of the population. Intuitively, the panel mirrors the structure of the population if when $x\%$ of the population has some particular characteristics, then $x\%$ of the panel consists of individuals with these characteristics. For example, if the female share of the population is $48\%$, then $48\%$ of the panel should be females. We call this intuitive property as *proportional representation*. In the example below, we show that measuring representation through $q$-social cost may fail to achieve proportional representation.

| | $A$ | $B$ | $C$ | D |
|---|---|---|---|---|
| $A$ | 0 | $\infty$ | $\infty$ | $\infty$ |
| $B$ | $\infty$ | 0 | $\infty$ | $\infty$ |
| $C$ | $\infty$ | $\infty$ | 0 | 10 |
| $D$ | $\infty$ | $\infty$ | 10 | 0 |

**Example 1.** *Let $n$ be odd, $k = 3$ and $q = 1$. Assume that there are four group of individuals, $A$, $B$, $C$ and $D$. There are exactly one individual in group $A$, and exactly one individual in group $B$, while there are $\frac{n-1}{2}$ individuals in group $C$ and $\frac{n-1}{2}$ individuals in group $D$. The distances between individuals in different groups is specified in the following table.*
*It is not difficult to see that any panel with minimum social cost contains the single individuals in groups $A$ and $B$ and one individual from either group $C$ or group $D$, as otherwise the social cost would be unbounded. This means that while the individuals in group $C$ form almost $50\%$ of the population, and similarly do the individuals in group $D$, in any panel with optimal social cost, either group $C$ or $D$ is not represented at all. On the other hand, the two eccentric individuals are always part of the panel.*

Uniform selection seems to achieve proportional representation *ex ante* (before the randomness is realized), since in expectation the selected panel reflects the composition of the population, especially when the size of the panel is very large. However, practitioners are often not satisfied with the possibility of having a realized panel that does not reflect the demographic characteristics of the population over a predefined set of features. For example, if the population is split evenly between college-educated and college-uneducated individuals, it is still possible for uniform selection to return a panel that consists only of college-educated individuals. Therefore, practitioners demand selection procedures that satisfy proportional representation *ex post* (after the randomness is realized). The question arises: *Under what conditions can a panel be considered proportionally representative of a population in a rigorous way*?

One approach is to consider a panel proportionally representative when each individual feature is proportionally represented. For instance, if $49\%$ of the population is college-educated and $51\%$ of the population is college-uneducated, a panel is feature-wise proportionally representative, if $49\%$ of the representatives are college-educated and $51\%$ of the representatives are college-uneducated. However, only requiring proportional representation of individual features may result in panels that do not accurately reflect the population as a whole. For example, when a population consists of $49\%$ college-educated high-income individuals and $49\%$ college-uneducated low-income individuals, a panel in which half the people are college-educated low-income and the other half are college-uneducated high-income may be incorrectly perceived as proportionally representative. One way to prevent this issue is to consider all the intersectional groups, and characterize a panel as representative when each group is proportionally represented. Nevertheless, in cases where the predefined set of features is large, the number of possible groups can grow exponentially, making it infeasible to represent all of them.

In this section, we borrow a notion of proportional representation by recent works on multiwinner elections and fair allocation of public goods Aziz et al. [2017], Fain et al. [2018], Conitzer et al. [2019], Cheng et al. [2020], called the *core*. The main idea of the core is: *Every subset $S$ of the population is entitled to choose up to $|S|/n \cdot k$ representatives*. Formally, a panel $P$ is called proportionally representative, or is said to be in the core, if there does not exist a subset $S$ of the population that could choose a panel $P'$, with $|P'| \leqslant |S|/n \cdot k$, under which all of them feel more represented. Note that this notion is not defined over predefined groups using particular features, but it provides fair representation

in the panel to any subset of the population.

Before we formally define core, we first introduce the following definition: For $\alpha \geqslant 1$, the $\alpha$-$q$-*preference count* of $P$ with respect to $P'$ is the number of individuals whose $q$-cost under $P$ is larger than $\alpha$ times their $q$-cost under $P'$:

$$V_q(P, P', \alpha) = |\{i \in [n] : c_q(i, P) > \alpha \cdot c_q(i, P')\}|.$$

A panel $P$ is *in the* $\alpha$-$q$-*core*, if for any panel $P'$, $V_q(P, P', \alpha) < |P'| \cdot n/k$. For $\alpha = 1$, we say that the panel is in the $q$-core. We define $\alpha$-$q$-core for $\alpha > 1$, since even when $q = 1$, a panel in the exact $q$-core is not guaranteed to exist [Chen et al., 2019, Micha and Shah, 2020] .

*Ex Post $q$-Core.* A selection algorithm $\mathsf{A}_{k,q}$ is in the *ex post* $\alpha$-$q$-*core* (or ex post $q$-core, for $\alpha = 1$) if every panel $P$ in the support of $\mathsf{A}_{k,q}$ is in the $\alpha$-$q$-core, i.e., for all $P$ drawn from $\mathsf{A}_{k,q}$ and all $P'$,

$$V_q(P, P', \alpha) < |P'| \cdot n/k.$$

*Ex Ante $q$-Core.* A selection algorithm $\mathsf{A}_{k,q}$ is in the *ex ante* $\alpha$-$q$-*core* (or ex ante $q$-core, for $\alpha = 1$) if for all $P'$:

$$\mathsf{E}_{P \sim \mathsf{A}_{k,q}}[V_q(P, P', \alpha)] < |P'| \cdot \frac{n}{k}.$$

The idea of requiring a core-like property over the expected number of preference counts was introduced by Cheng et al. [2020] in a mutliwinner election setting.

### 4.4.1 Fairness and The Ex Post Core

Here, we investigate if there are selection algorithms that are fair and provide a good approximation to the ex post $q$-core.

We start by showing that when $q = k$, uniform selection is in the ex post 2-$k$-core. This positive result derives from the fact that a given panel is not in the $k$-core if and only if all individuals agree that there exists another panel of size $k$ that is more representative for all of them.

**Theorem 13.** *For $q = k$, uniform selection is in the ex post 2-core and this is tight.*

*Proof.* Consider a panel $P$ returned by uniform selection. It suffices to show that for any arbitrary panel $P'$ of size $k$, the $q$-cost of all individuals cannot be improved by a factor of greater than $\alpha = 2$.

Let $i_1$ and $i_2$ be the two individuals in the population with the maximum distance between them. Now, consider an arbitrary representative $r$ in panel $P'$. Without loss of generality, suppose that $c_k(i_1, P') \leqslant c_k(i_2, P')$. Then, we have

$$\begin{aligned} c_k(i_2, P) = \max_{j \in P} d(i_2, j) &\leqslant d(i_1, i_2) && \text{(by the choice of } i_1 \text{ and } i_2\text{)} \\ &\leqslant d(i_1, r) + d(r, i_2) && \text{(triangle inequality)} \\ &\leqslant c_k(i_1, P') + c_k(i_2, P') && \text{(as } r \in P'\text{)} \\ &\leqslant 2 \cdot c_k(i_2, P'). \end{aligned}$$

This implies that for any panel $P$ in $\mathcal{U}_k$, $V_q(P, P', 2) < |P'| \cdot n/k = n$, since $q$-cost for $i_2$ does not improve by a factor of more than two. Thus, uniform selection is in the ex post 2-core.

Next, we show that there exists an instance such that uniform selection is not in the ex post $\alpha$-$k$-core for $\alpha < 2$. Consider the case that the individuals are assigned into three groups, $A$, $B$ and $C$, with $\lfloor k/2 \rfloor$, $\lceil k/2 \rceil$, and $n - k$ individuals, respectively. The distances between individuals is as specified in the following table.

| | $A$ | $B$ | $C$ |
|---|---|---|---|
| $A$ | 0 | 2 | 1 |
| $B$ | 2 | 0 | 1 |
| $C$ | 1 | 1 | 0 |

Panel $P$ consisting of all the $k$ people in groups $A$ and $B$ is in the support of uniform selection. Then, for $i \in A \cup B$, $c_k(i, P) = 2$ as the $k$-th closest representative in $P$ lies in the other group. For $i \in C$, the $c_k(i, P) = 1$. Now, consider panel $P'$ that consists of $k$ individuals from group $C$. The $q$-costs of all individuals improve by a factor of at least 2. Hence, $\mathcal{U}_k$ violates ex post 2-$k$-core in this example. □

Note that the above theorem shows that any selection algorithm is in the ex post 2-core when $q = k$, since the only property of uniform selection that is utilized is that it returns a panel of size $k$.

Next, we show that for any value of $q$ other than $k$, uniform selection does not satisfy any bounded approximation of the ex post $q$-core.

**Theorem 14.** *For any $q \in [k-1]$ and $\lfloor n/k \rfloor \geqslant k$, there exists an instance such that uniform selection is not in the ex post $\alpha$-$q$-core for any bounded $\alpha$.*

*Proof.* Consider an instance in which there are $\lfloor n/k \rfloor$ individuals are in group $A$ and the remaining individuals in group $B$. Suppose that the distance between two individuals in the same group is 0, and the distance between two individuals in different groups is 1. Since, $\lfloor n/k \rfloor \geqslant k$, uniform selection has a non-zero probability of returning a panel where all the representatives are from group $A$. In this scenario, for any $q \in [k-1]$, the $q$-cost of all the individuals in group $B$ is equal to 1. However, individuals in group $B$ are entitled to choose up to $k-1$ representatives among themselves, and if they do so, their $q$-cost becomes 0, resulting in an unbounded improvement in cost. Therefore, uniform selection is not in the ex post $\alpha$-$q$-core for any bounded $\alpha$. □

Uniform selection fails to provide any bounded approximation to the ex post $q$-core, for nearly all $q$, except for $q = k$ where it provides an approximation of 2. Therefore, we ask: For any given $q$, is there any selection algorithm that keeps the fairness guarantee of uniform selection (i.e. fairness) and ensures that every panel it returns is in the approximate $q$-core? We answer this positively.

We present a selection algorithm, called FAIRGREEDYCAPTURE$_{k,q}$, that is fair and in the ex post $((5+\sqrt{41})/2)$-$q$-core. Our algorithm leverages the basic idea of the Greedy Capture algorithm introduced by Chen et al. [2019], which returns a panel in the approximate 1-core. Briefly, Greedy Capture starts with an empty panel and grows a ball around any individual at the same rate. When a ball captures at least $\lceil n/k \rceil$ individuals for the first time, the center of the ball is included in the panel and all the captured individuals are disregarded. As the balls continue to grow, when new individuals are captured by balls whose center is in the panel, they are disregarded.

Similarly, FAIRGREEDYCAPTURE$_{k,q}$, in Algorithm 3, starts with an empty panel $P$ and grows a ball around every individual in $[n]$ at the same rate. When a ball captures $\lceil q \cdot n/k \rceil$ individuals (if more than $\lceil q \cdot n/k \rceil$ individuals have been captured, it chooses exactly $\lceil q \cdot n/k \rceil$ by arbitrarily excluding some points on the boundary), the algorithm selects $q$ of them uniformly at random, includes them in the panel $P$, and disregards all the $\lceil q \cdot n/k \rceil$ individuals. When this happens, we say that the algorithm *detects* this ball. Unlike Greedy Capture, we continue growing balls only around the individuals that are not yet disregarded, i.e. detected balls are frozen. When fewer than $\lceil q \cdot n/k \rceil$ individuals remain, the algorithm

**Algorithm 3** FAIRGREEDYCAPTURE$_{k,q}$

**Input:** $[n]$, $d$
**Output:** Panel $P$
$R \leftarrow [n]; \delta \leftarrow 0; P \leftarrow \emptyset$
**while** $|R| \geqslant \lceil q \cdot n/k \rceil$ **do**
  Smoothly increase $\delta$
  **while** $\exists i \in R$ such that $|B(i,\delta) \cap R| \geqslant \lceil q \cdot n/k \rceil$ **do**
    $S \leftarrow \lceil q \cdot n/k \rceil$ individuals arbitrary chosen from $B(i,\delta)$
    $\widehat{P} \leftarrow$ pick $q$ individuals from $S$ uniformly at random
    $P \leftarrow P \cup \widehat{P}$
    $R \leftarrow R \setminus S$
  **end while**
**end while**
$\widehat{P} \leftarrow k - |P|$ individuals from $[n] \setminus P$ by picking $i \in R$ with probability $k/n$ and $i \in [n] \setminus (P \cup R)$ with probability $\frac{k-|P|-|R|\cdot k/n}{n-|P|-|R|}$
$P \leftarrow P \cup \widehat{P}$

selects the remaining representatives from among the individuals who have not yet been included in the panel as follows: each individual who has not been disregarded is selected with probability $k/n$, and the remaining probability mass is allocated uniformly among the individuals who have been disregarded but not selected. This can be achieved through systematic sampling Yates [1948].

To prove that FAIRGREEDYCAPTURE$_{k,q}$ is fair and in the approximate ex post $q$-core, we need the next two technical lemmas.

**Lemma 12.** *Let $S \subseteq [n]$ and $P'$ be a panel.*

1. *There exists a partition $T_1, \ldots, T_m$ of $S$ with respect to $P'$, with $m \leqslant \lfloor |P'|/q \rfloor$, such that for each $\ell \in [m]$, there is $i^*_\ell \in T_\ell$ for which it holds that $c_q(i, P') \leqslant c_q(i^*_\ell, P')$ and $\mathsf{top}_q(i, P') \cap \mathsf{top}_q(i^*_\ell, P') \neq \emptyset$, for any $i \in T_\ell$.*

2. *There exists a partition $T_1, \ldots, T_m$ of $S$ with respect to $P'$, with $m \leqslant \lfloor |P'|/q \rfloor$, such that for each $\ell \in [m]$, there is $i^*_\ell \in T_\ell$ for which it holds that $c_q(i^*_\ell, P') \leqslant c_q(i, P')$ and $\mathsf{top}_q(i, P') \cap \mathsf{top}_q(i^*_\ell, P') \neq \emptyset$, for any $i \in T_\ell$.*

*Proof.* We start by showing the first part. We partition all the individuals in $S$into $m \leqslant \lfloor |P'|/q \rfloor$ groups, denoted by $T_1, \ldots, T_m$ iteratively as follows.

Suppose $i^*_1$ is the individual with the smallest $q$-cost over $P'$ (ties are broken arbitrary), i.e. $i^*_1 = \arg\max_{i \in S} c_q(i, P')$. Then, $T_1$ is the set of all the individuals whose $q$ closest representatives from $P'$ includes at least one member of $\mathsf{top}_q(i^*_1, P')$, i.e.

$$T_1 = \{i \in S : \mathsf{top}_q(i, P') \cap \mathsf{top}_q(i^*_1, P') \neq \emptyset\}.$$

Next, from the remaining individuals, suppose $i^*_2$ is the one with the smallest $q$-cost over $P'$, i.e. $i^*_2 = \arg\min_{i \in S \setminus T_1} c_q(i, P)$. Construct $T_2$ from $S \setminus T_1$ similarly by taking all the individuals whose at least one of their $q$ closest representatives in $P'$ is included in $\mathsf{top}_q(i^*_2, P')$. We repeat this procedure, and

in round $\ell$, we find $i^*_\ell \in S \setminus (\cup_{\ell'=1}^{\ell-1} T_{\ell'})$ that has the smallest cost over $P'$, and construct $T_\ell$ by assigning any individual in $S \setminus (\cup_{\ell'=1}^{\ell-1} T_{\ell'})$ whose at least one of the $q$ closest representatives belongs in $\mathsf{top}_q(i^*_\ell, P')$. Note that for any $\ell_1, \ell_2 \in [m]$ with $\ell_1 < \ell_2$, $\mathsf{top}_q(i^*_{\ell_1}, P') \cap \mathsf{top}_q(i^*_{\ell_2}, P') = \emptyset$, as if at least one of the $q$ closest representatives of $i^*_{\ell_2}$ in $P$ is included in $\mathsf{top}_q(i^*_{\ell_1}, P')$, then $i^*_{\ell_2}$ would have been assigned to $T_{\ell_1}$ and would not belong in $S \setminus (\cup_{\ell'=1}^{\ell_2-1} T_{\ell'})$. This means that in each round, we consider $q$ representatives that have not been considered before, and hence after $\lfloor |P'|/q \rfloor$ rounds, less than $q$ representatives in $P'$ may remain unconsidered. As a result, after at most $\lfloor |P'|/q \rfloor$ rounds, all the individuals will have been assigned to some group, since at least one of their $q$ closest representatives has been considered.

The second part follows by simply setting $i^*_\ell$ to be equal to the individual in $S \setminus (\cup_{\ell'=1}^{\ell-1} T_{\ell'})$ that has the largest cost over $P'$, i.e. $i^*_\ell = \arg\min_{i \in S \setminus (\cup_{\ell'=1}^{\ell-1} T_{\ell'})} c_q(i, P)$. All the remaining arguments remain the same. □

**Lemma 13.** *For any panel $P$ and any $i, i' \in [n]$, it holds that $c_q(i, P) \leqslant d(i, i') + c_q(i', P)$.*

*Proof.* Consider a ball centered at $i'$ with radius $c_q(i', P)$. This ball contains at least $q$ representatives of $P$. Hence, $c_q(i, P)$ is less than or equal to the distance of $i$ to one of the $q$ representatives that are included in $B(i', c_q(i', P))$ which is at most $d(i, i') + c_q(i', P)$. □

Now, we are ready to prove the next theorem.

**Theorem 15.** *For any $q$,* FAIRGREEDYCAPTURE$_{k,q}$ *is fair and in the ex post $((5 + \sqrt{41})/2)$-$q$-core.*

*Proof.* We start by proving the fairness of FAIRGREEDYCAPTURE$_{k,q}$.

**Lemma 14.** FAIRGREEDYCAPTURE$_{k,q}$ *is fair.*

*Proof.* Suppose that $q \cdot n/k$ is an integer. Then, each individual that is disergarded in the while loop of the algorithm is included in the panel with probability exactly $k/n$. Now, suppose that after the algorithm has detected $t$ balls, less than $q \cdot n/k$ individuals have not been disregarded. Then, when the algorithm exits the while loop we have that $|R| = n - t \cdot q \cdot n/k$ and $k - |P| = k - t \cdot q$. But since,

$$|R| \cdot k/n = k - t \cdot q,$$

we conclude that the remaining $k - t \cdot q$ representatives are chosen uniformly among the individuals in $R$. Thus, the algorithm returns a panel of size $k$ and each $i \in [n]$ is chosen with probability $k/n$.

Now, we focus on the case where $q \cdot n/k$ is not an integer. In this case, note that in the while loop of the algorithm, less than $k$ representatives are included in the panel, since $q$ representatives are included in it every time that $\lceil q \cdot n/k \rceil$ non-disregarded individuals are captured from a ball. Moreover, each individual that is disregarded is chosen with probability strictly less than $k/n$. Now suppose that after exiting the while loop, there are individuals that have not been disregarded, i.e. $|R| > 0$. First, we show that the algorithm correctly chooses another $k - |P|$ representatives and outputs a panel of size $k$. The algorithm would select each individual in $R$ with probability $k/n$ and allocates the remaining probability — which is equal to $k - |P| - |R| \cdot k/n$ —uniformly among the $n - |P| - |R|$ individuals that have been disregarded but not selected in $P$. To satisfy fairness for people in $R$, it suffices to show that $|R| \cdot k/n < k - |P|$. Since for each individual $i \in [n] \setminus R$ we have $\Pr[i \in P] = q/\lceil q \cdot n/k \rceil < k/n$, then we have $|P| = \mathsf{E}[|P|] = \sum_{i \in [n] \setminus R} \Pr[i \in P] < (n - |R|) \cdot k/n$. Thus, $k - |P| > k - (n - |R|) \cdot k/n = |R| \cdot k/n$. Hence, the algorithm outputs panels of size $k$.

It remains to show that each individual in $[n] \setminus R$, which is disregarded in the while loop, is included in the panel with probability $k/n$. First, note that all of them are included in the panel with the same probability. This holds, since each is selected with probability $q/\lceil q \cdot n/k \rceil$ from the ball that captured them in the while loop, and, when not selected in the while loop, they get an equal chance of selection of $\frac{k-|P|-|R|\cdot k/n}{n-|R|-|P|}$. Since the size of the final panel returned by the algorithm is always $k$, and by linearity of expectation, we have $k = |R| \cdot k/n + \sum_{i\in[n]\setminus R} \Pr[i]$. By equality of $\Pr[i]$'s, we conclude that all must be equal to $k/n$ and each individual in $[n]$ is included in the panel with probability $k/n$. □

We proceed by showing that FAIRGREEDYCAPTURE$_{k,q}$ is in the ex post $((5+\sqrt{41})/2)$-$q$-core. Let $P$ be any panel that the algorithm may return. Suppose for contradiction that there exists a panel $P'$ such that $V_q(P, P', (5+\sqrt{41})/2) \geqslant |P'| \cdot n/k$. This means that there exists $S \subseteq [n]$, with $|S| \geqslant |P'| \cdot n/k$, such that:

$$\forall i \in S, \qquad c_q(i, P) > (5+\sqrt{41})/2 \cdot c_q(i, P'). \tag{4.3}$$

Let $T_1, \ldots, T_m$ be a partition of $S$ with respect to $P'$, as given in the first part of Lemma 12. Since $m \leqslant \lfloor |P'|/q \rfloor$ and $|S| \geqslant |P'| \cdot n/k$, we conclude that there exists a part, say $T_\ell$, that has size at least $q \cdot n/k$. Since, there exists $i^*_\ell \in T_\ell$ such that for each $i \in T_\ell$ it holds that $c_q(i, P') \leqslant c_q(i^*_\ell, P')$ and $\mathsf{top}_q(i, P') \cap \mathsf{top}_q(i^*_\ell, P') \neq \emptyset$, we can conclude that $d(i^*_\ell, i) \leqslant 2 \cdot c_q(i^*_\ell, P')$, as following: Pick an arbitrary representative in $\mathsf{top}_q(i, P') \cap \mathsf{top}_q(i^*_\ell, P')$ and denote it as $r_i$. Then,

$$d(i, i^*_\ell) \leqslant d(i, r_i) + d(r_i, i^*_\ell) \leqslant c_q(i, P') + c_q(i^*_\ell, P') \leqslant 2 \cdot c_q(i^*_\ell, P').$$

This means that there exists a ball centered at $i^*_\ell$ that has radius $2 \cdot c_q(i^*_\ell, P')$ and captures all the individuals in $T_\ell$. Since there are sufficiently many individuals in this ball, it is possible that the algorithm detected this ball (or a nested one) during its execution. If this happened, this means that there are $q$ representatives in $P$ that are located within the ball $B(i^*_\ell, 2 \cdot c_q(i^*_\ell, P'))$. Then, we get that $c_q(i^*_\ell, P) \leqslant 2 \cdot c_q(i^*_\ell, P')$ which contradicts Equation (4.3).

On the other hand, if the algorithm did not detect this ball (or a nested one) during its execution, this means that some of the individuals in $T_\ell$ have been disregarded before the ball centered at $i^*_\ell$ captures sufficiently many of them. Hence, some individuals in $T_\ell$ have been captured from a different ball with radius at most $2 \cdot c_q(i^*_\ell, P')$. Suppose that $i'$ is the first individual in $T_\ell$ that was captured from such a ball. Then, we get that $q$ representatives in $P$ are within this ball, which means that $c_q(i', P) \leqslant 4 \cdot c_q(i^*_\ell, P')$, since the distance of $i'$ form any individual in this ball is at most equal to the diameter of the ball. We consider the minimum multiplicative improvement of both $i^*_\ell$ and $i'$:

$$\begin{aligned}
&\min\left(\frac{c_q(i', P)}{c_q(i', P')}, \frac{c_q(i^*_\ell, P)}{c_q(i^*_\ell, P')}\right) \\
\leqslant &\min\left(\frac{c_q(i', P)}{c_q(i', P')}, \frac{d(i^*_\ell, i') + c_q(i', P)}{c_q(i^*_\ell, P')}\right) && \text{(by Lemma 13)} \\
\leqslant &\min\left(\frac{c_q(i', P)}{c_q(i', P')}, \frac{d(i^*_\ell, r_{i'}) + d(i', r_{i'}) + c_q(i', P)}{c_q(i^*_\ell, P')}\right) && \text{(by Triangle Inequality)} \\
\leqslant &\min\left(\frac{c_q(i', P)}{c_q(i', P')}, \frac{c_q(i^*_\ell, P') + c_q(i', P') + c_q(i', P)}{c_q(i^*_\ell, P')}\right) && \text{(as } r_{i'} \in \mathsf{top}_q(i', P') \cap \mathsf{top}_q(i^*_\ell, P')\text{)} \\
\leqslant &\min\left(\frac{4 \cdot c_q(i^*_\ell, P')}{c_q(i', P')}, 5 + \frac{c_q(i', P')}{c_q(i^*_\ell, P')}\right) && \text{(as } c_q(i', P) \leqslant 4 \cdot c_q(i^*_\ell, P')\text{)} \\
\leqslant &\max_{z \geqslant 0} \min(4 \cdot z, 5 + 1/z) = (5+\sqrt{41})/2
\end{aligned}$$

which violates Equation (4.3) and the theorem follows. □

In the next section, we show that no fair algorithm provides an approximation better than 2 to the ex ante $q$-cost, for any $q$. Since, the ex post $\alpha$-$q$-core implies the ex ante $\alpha$-$q$-core, we get that no fair selection algorithm provides an approximation better than 2 to the ex post $q$-core, and therefore FAIRGREEDYCAPTURE$_{k,q}$ is optimal up to a small constant factor.

### 4.4.2 Uniform Selection and the Ex Ante Core

We have already discussed that uniform selection fails to provide any reasonable approximation to the ex post $q$-core, for almost all values of $q$. However, as we mentioned in the introduction, it seems to satisfy the ex ante $q$-core, at least when $k$ is very large. In this section, we ask whether indeed uniform selection satisfies the ex ante $q$-core, in a rigorous way, for all values of $q$ and $k$. We show that uniform selection is in the 4-ex post $q$-core, for any $q$.

To show this result, we use the following form of Chu–Vandermonde identity which we prove for completeness.

**Definition 7** (Chu–Vandermonde identity)**.** *For any $n, k,$ and $r$, with $0 \leqslant r \leqslant k \leqslant n$, it holds*

$$\sum_{j=0}^{n} \binom{j}{r} \cdot \binom{n-j}{k-r} = \binom{n+1}{k+1}.$$

*Proof.* We give a combinatorial argument for this identity. Suppose we want to select $k+1$ items out of a set of size $n+1$. For $i \in [1, n+1]$, let $P_i$ be the number of such subsets in which the $(r+1)$th picked item is item $i$. As each subset is counted exactly once among $P_i$'s (at the position of its $(r+1)$th item), we have $\sum_{i=1}^{n+1} P_i = \binom{n+1}{k+1}$. Now, we calculate $P_i$. Suppose the $(r+1)$th item is $i$. Then, $r$ items should be selected from the first $i-1$ items and $k+1-(r+1) = k-r$ items should be selected from the last $n+1-i$ items. Therefore, $P_i = \binom{i-1}{r} \cdot \binom{n-(i-1)}{k-r}$. Then, we have

$$\binom{n+1}{k+1} = \sum_{i=1}^{n+1} P_i = \sum_{i=1}^{n+1} \binom{i-1}{r} \cdot \binom{n-(i-1)}{k-r} = \sum_{j=0}^{n} \binom{j}{r} \cdot \binom{n-j}{k-r}.$$ □

Now, we are ready to prove the following theorem.

**Theorem 16.** *For any $q$, uniform selection is in ex ante $4$-$q$-core, i.e. for any panel $P'$*

$$\mathsf{E}_{P \sim \mathcal{U}_k}[V_q(P, P', 4)] < |P'| \cdot \frac{n}{k}.$$

*Proof.* Let $P'$ be any panel. By linearity of expectation, we have that

$$\mathsf{E}_{P \sim \mathcal{U}_k}[V_q(P, P', 4)] = \sum_{i \in [n]} \Pr_{P \sim \mathcal{U}_k}[c_q(i, P) > 4 \cdot c_q(i, P')].$$

Let $T_1 \ldots, T_m$ be a partition of $[n]$ with respect to $P'$, as given in the second part of Lemma 12. For each $\ell \in [m]$, we reorder the individuals in $T_\ell$ in an increasing order based on their distance from $i^*_\ell$, and relabel them as $i^\ell_1, \ldots, i^\ell_{|T_\ell|}$. This way, $i^\ell_1$ and $i^\ell_{|T_\ell|}$ are the individuals in $T_\ell$ that have the smallest and the largest distance from $i^*_\ell$, respectively. Then, we get that

$$\sum_{i \in [n]} \Pr_{P \sim \mathcal{U}_k}[c_q(i, P) > 4 \cdot c_q(i, P')] = \sum_{\ell=1}^{m} \sum_{j=1}^{|T_\ell|} \Pr_{P \sim \mathcal{U}_k}[c_q(i^\ell_j, P) > 4 \cdot c_q(i^\ell_j, P')]. \tag{4.4}$$

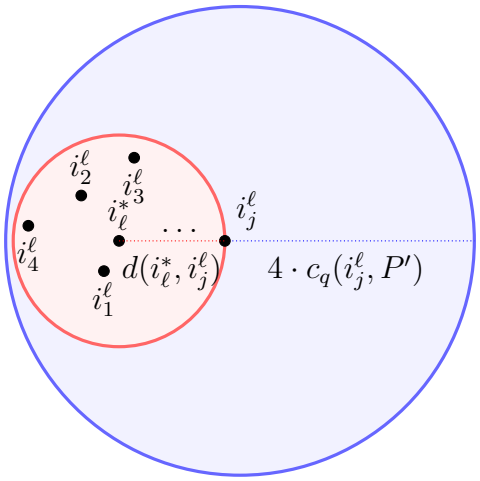


Figure 4.1: Diagram for Proof of Lemma 15

In the next lemma, we bound $\Pr_{P\sim\mathcal{U}_k}[c_q(i_j^\ell, P) > 4 \cdot c_q(i_j^\ell, P')]$ for each $i_j^\ell$.

**Lemma 15.** *For each $\ell \in [m]$ and $j \in [|T_\ell|]$, $\Pr_{P\sim\mathcal{U}_k}[c_q(i_j^\ell, P) > 4 \cdot c_q(i_j^\ell, P')] \leqslant \sum_{r=0}^{q-1} \frac{\binom{j}{r}\binom{n-j}{k-r}}{\binom{n}{k}}$.*

*Proof.* For each $i_j^\ell$, let $r_j^\ell$ be an arbitrary representative in $\mathsf{top}_q(i_j^\ell, P') \cap \mathsf{top}_q(i_\ell^*, P')$. Then, we get that

$$d(i_j^\ell, i_\ell^*) \leqslant d(i_j^\ell, r_j^\ell) + d(r_j^\ell, i_\ell^*) \leqslant c_q(i_j^\ell, P') + c_q(i_\ell^*, P') \leqslant 2 \cdot c_q(i_j^\ell, P'), \tag{4.5}$$

where the last inequality follows from the fact that $i_\ell^*$ has the smallest cost over $P'$ among all the individuals in $T_\ell$. Now, consider the ball that is centered at $i_j^\ell$ and has radius $4 \cdot c_q(i_j^\ell, P')$. Note that this ball contains any individual $i_{j'}^\ell$ with $j' < j$. Indeed, for each $i_j^\ell$ and $i_{j'}^\ell$ with $j' < j$, we have that

$$d(i_j^\ell, i_{j'}^\ell) \leqslant d(i_j^\ell, i_\ell^*) + d(i_\ell^*, i_{j'}^\ell) \leqslant 2 \cdot d(i_j^\ell, i_\ell^*) \leqslant 4 \cdot c_q(i_j^\ell, P'),$$

where the second inequality follows form the fact that for each $j', j \in [|T_\ell|]$ with $j' < j$, $d(i_{j'}^\ell, i_\ell^*) \leqslant d(i_j^\ell, i_\ell^*)$ and the last inequality follows form Equation (4.5). This argument is drawn in Figure 4.1.

When $c_q(i_j^\ell, P) > 4 \cdot c_q(i_j^\ell, P')$, then we get that $|P \cap \{i_1^\ell, \ldots, i_j^\ell\}| < q$, as otherwise there would exist at least $q$ individuals in $B(i_j^\ell, 4 \cdot c_q(i_j^\ell, P'))$, and $c_q(i_j^\ell, P)$ would be at most $4 \cdot c_q(i_j^\ell, P')$. Hence, we have that

$$\begin{aligned}
\Pr_{P\sim\mathcal{U}_k}[c_q(i_j^\ell, P) > 4 \cdot c_q(i_j^\ell, P')] &\leqslant \Pr_{P\sim\mathcal{U}_k}\left[\,|P \cap \{i_1^\ell, \ldots, i_j^\ell\}| < q\right] \\
&= \Pr_{P\sim\mathcal{U}_k}\left[\bigcup_{r=0}^{q-1} |P \cap \{i_1^\ell, \ldots, i_j^\ell\}| = r\right] \\
&\leqslant \sum_{r=0}^{q-1} \Pr_{P\sim\mathcal{U}_k}\left[\,|P \cap \{i_1^\ell, \ldots, i_j^\ell\}| = r\right] = \sum_{r=0}^{q-1} \frac{\binom{j}{r}\binom{n-j}{k-r}}{\binom{n}{k}}
\end{aligned}$$

where the second inequality follows from the Union Bound and the last equality follows form the fact that uniform selection chooses $k$ out of $n$ individuals uniformly at random. □

Then, by returning to Equation (4.4) we get that,

$$\begin{aligned}
\mathsf{E}_{P\sim\mathcal{U}_k}[V_q(P,P',4)] &= \sum_{\ell=1}^{m}\sum_{j=1}^{|T_\ell|}\Pr_{P\sim\mathcal{U}_k}[c_q(i_j^\ell,P) > 4\cdot c_q(i_j,P')] \\
&\leqslant \sum_{\ell=1}^{m}\sum_{j=1}^{|T_\ell|}\sum_{r=0}^{q-1}\frac{\binom{j}{r}\binom{n-j}{k-r}}{\binom{n}{k}} = \sum_{\ell=1}^{m}\sum_{r=0}^{q-1}\sum_{j=1}^{|T_\ell|}\frac{\binom{j}{r}\binom{n-j}{k-r}}{\binom{n}{k}} \leqslant \sum_{\ell=1}^{m}\sum_{r=0}^{q-1}\sum_{j=0}^{n}\frac{\binom{j}{r}\binom{n-j}{k-r}}{\binom{n}{k}} \\
&\overset{(1)}{=} \sum_{\ell=1}^{m}\sum_{r=0}^{q-1}\frac{\binom{n+1}{k+1}}{\binom{n}{k}} = \sum_{\ell=1}^{m}\sum_{r=0}^{q-1}\frac{n+1}{k+1} = m\cdot q\cdot\frac{n+1}{k+1} < |P'|\cdot\frac{n}{k},
\end{aligned}$$

where (1) follows from Chu–Vandermonde identity and the last inequality follows from the facts that $m \leqslant \lfloor {}^{|P'|}/_{q} \rfloor$ and ${}^{(n+1)}/_{(k+1)} < {}^{n}/_{k}$, when $k < n$. □

In the next theorem, we show that for any $q < k$, no selection algorithm that is fair, is guaranteed to achieve ex ante $\alpha$-$q$-core with $\alpha < 2$, and hence uniform selection is optimal up to a factor of 2.

**Theorem 17.** *For any $q \in [k-1]$, when $n \geqslant 2k^2/(k-q)$, there exists an instance such that no selection algorithm that is fair, is in the ex ante $\alpha$-$q$-core with $\alpha < 2$.*

*Proof.* Consider a star graph with $n - q$ leaves and an internal node. Suppose $q$ individuals $I = \{i_1, \ldots, i_q\}$ lie on the internal node, and exactly one individual lies on each of the $n - q$ leaves. Individuals in $I$ have a distance of 0 from each other and a distance of 1 from $[n] \setminus I$; and, the distance between a pair of individuals from $[n] \setminus I$ is equal to 2. These distances satisfy the triangle inequality.

Let $P$ be an arbitrary panel of size $k$ that does not contain $i_1$. We show that for $P' = I$ and any $\alpha < 2$, we have that $V_q(P, P', \alpha) \geqslant n - k$. For any $i \in I$, it holds $c_q(i, P) = 1$ and $c_q(i, P') = 0$ — which is an unbounded improvement. For any individual $i$ in $[n] \setminus (I \cup P)$, $c_q(i, P) = 2$ since their $q$th closest representative in $P$ would be on another leaf, while $c_q(i, P') = 1$ — which yields a 2 factor improvement. Therefore, $V_q(P, P', \alpha) \geqslant |([n] \setminus (I \cup P)) \cup I| \geqslant n - |P| = n - k$.

Under any fair selection algorithm, $i_1$ is not included in the panel with probability $1 - {}^{k}/_{n}$. Thus, we have that

$$\mathsf{E}_{P\sim\mathcal{U}_k}[V_q(P,P',\alpha)] \geqslant \Pr[i_1 \notin P]\cdot(n-k) = (1-k/n)\cdot(n-k) \geqslant q\cdot n/k = |P'|\cdot n/k,$$

where the last inequality follows from the assumption that $n \geqslant 2k^2/(k-q)$. □

## 4.5 Discussion

In this chapter, we investigate whether panels that are selected at random reflect the composition of a population, when the invididulas lie in a metric space. We get two different approaches for measuring the representation of a panel in a rigorous way.

In Section 4.3, we use the $q$-social cost, for measuring the extend that a panel is representative. Our results in Section 4.3.2 show that in some cases, relaxing fairness requirements allows improving representation dramatically. More generally, it is interesting to understand the tradeoff between representation and fairness, and to chart the Pareto frontier.

In Section 4.4, we show that the $q$-social cost fails to capture the basic property of proportional representation. Thus, we turn our attention to the idea of the $q$-core of measuring the proportional representation of a panel. There are many interesting directions for future work from this approach as well. We show that FAIRGREEDYCAPTURE is in the $((5+\sqrt{41})/2)$-$q$-core, but we do not provide a lower bound that indicates if this analysis is tight. Finding tight bounds for the general case is an open question. Moreover, we show that there is no fair algorithm that is in the ex post $\alpha$-$q$-core for $\alpha < 2$. Closing the gap between this lower bound and the bound that FAIRGREEDYCAPTURE guarantees is another interesting direction. A limitation of our algorithm is that it is parameterized by $q$, since in some applications, it may be difficult to choose the appropriate value of $q$. Exploring whether we can achieve good guarantees without knowing $q$ is interesting future direction. Moreover, as we mentioned above, solely focusing on proportional representation by individual features could result in panels that are not genuinely representative of the population. However, it is important to note that while a panel in the core guarantees proportional representation over intersectional features, it may not necessarily ensure proportional representation over individual features. We defer the exploration of simultaneously achieving proportional representation over both intersectional and individual features to future research.

In both sections, we focused our attention on the $q$-cost formulation of Caragiannis et al. [2022], in which each individual measures their distance to the $q$-th closest panel member. One can consider different cost functions. For example, what if different individuals use different values of $q$? Another appealing choice is when each individual measures the *average* distance to all panel member.

Lastly, it is interesting to explore in which cases the core and the social cost are incompatible, and investigate if there are algorithms that strike a balance between the two different objectives.

We hope that answers to some of these questions will lead to a better understanding of the strengths of sortition, and to new ways of realizing this democratic paradigm. Overall, we view the need for in-depth rigorous analyses of sortition algorithms as an urgent step in using this renewed paradigm to improve our democratic systems.

# Chapter 5

# Individual Fairness in Participatory Budgeting

## 5.1 Introduction

In this chapter, we consider the problem of selecting a fixed-size subset of candidates (a so-called *committee*) based on the approval preferences of voters. This problem has been extensively studied in recent years [Lackner and Skowron, 2023] and has a wide variety of applications, including political elections [Brill et al., 2020], recommender systems [Skowron et al., 2017], medical diagnostic decision-making [Gangl et al., 2019], and participatory budgeting [Peters et al., 2021b].

A central concern in committee voting is the principle of *proportional representation*, which states that the voters' interests and opinions should be reflected proportionately in the committee. While proportional representation is intuitive to understand in scenarios such as apportioning parliamentary seats based on vote shares [Balinski and Young, 1982, Pukelsheim, 2014], it is less straightforward to formalize in the context of approval-based committee elections. Indeed, the literature has defined a number of different concepts aiming to capture proportional representation.

Most (if not all) of these approval-based proportionality notions focus on the representation of *groups* of voters. Specifically, it is usually required that each sufficiently large group of voters is "represented" in the committee,[1] where the interpretation of "representation" differs across different notions. For example, *extended justified representation* [Aziz et al., 2017] prescribes that there exists at least one voter in the group approving a certain number of committee members, whereas *proportional justified representation* [Sánchez-Fernández et al., 2017] demands that there are sufficiently many committee members that are each approved by at least one voter in the group. Notably, neither definition comprises any representation requirements for *individual* voters in a group. Thus, a group may count as "represented" even though some voters in the group do not approve a single committee member.

In this chapter, we adopt an *individualistic* point of view: our goal is to provide all members of a voter group equal guarantees. Intuitively, when a population consists of $n$ voters and a committee of $k$ representatives is elected, we expect every cohesive voter group of size $\ell \cdot n/k$ to be represented by $\ell$ representatives in the committee; thus, each individual group member might reasonably hope that at least $\ell$ candidates represent her in the committee. This notion, which we call *individual representation*, is aligned with the notion of "individual fairness" that was recently introduced in clustering (and in

[1] Often, there is also a condition on the "cohesiveness" of the group, stating that the approval preferences of group members need to be sufficiently aligned.

particular in facility location problems) by Jung et al. [2020]: there, each individual expects to be served by a facility in distance proportional to the radius of the ball that captures its $n/k$ closest neighbors, where $n$ is the number of individuals and $k$ is the number of facilities.

Individual representation, as defined in this chapter, is a strengthening of a notion called *semi-strong justified representation* by Aziz et al. [2017]. The latter property requires that all members of a group are represented in the committee *at least once*, given that the group is large and cohesive enough. Individual representation strengthens this requirement by demanding that all members of cohesive groups are represented multiple times (in proportion to the group size). Aziz et al. [2017] observed that semi-strong JR cannot be provided in all instances; this immediately implies our stronger requirement is not universally attainable either.

In this chapter, we systematically study *individual representation (IR)*. Notwithstanding the observation that IR demands cannot always be met, we clarify how IR relates to existing axioms and we show that a large range of common approval-based committee voting rules can fail to provide IR even in cases where IR is achievable. We observe that even committees *approximating* IR may fail to exist. Moreover, we answer a question by Aziz et al. [2017] by showing that it is computationally intractable to decide whether a given instance admits a committee providing semi-strong JR or individual representation. We then turn our attention to *restricted domains* of preferences [Elkind and Lackner, 2015, Yang, 2019] and demonstrate that positive results can be obtained. Doing so, we uncover a striking difference between the *candidate interval* and *voter interval* domains: whereas the former restriction does not admit any non-trivial approximation of IR, we devise an efficient algorithm for selecting committees approximating IR for the latter. This is surprising insofar as these two domain restrictions often exhibit similar behavior [Pierczyński and Skowron, 2022, Terzopoulou et al., 2021]. Finally, we experimentally study how often IR is achievable for a wide variety of generated preference data, and how often established voting rules select IR outcomes.

## 5.2 Preliminaries

For $t \in \mathbb{N} = \{1, 2, \dots\}$, let $[t]$ denote the set $\{1, 2, \dots, t\}$.

Let $N = [n]$ be a set of $n$ *voters* (or *agents*) and $C = \{c_1, \dots, c_m\}$ be a set of $m$ *candidates*. Each voter $i \in N$ approves a subset $A_i \subseteq C$ of candidates. An *(approval) profile* $A = (A_1, \dots, A_n)$ contains the approval set $A_i$ for each $i \in N$.

Given a committee size $k \in [m]$, we want to select a subset $W \subseteq C$ of size $|W| = k$, referred to as a *committee*. We call $(A, k)$ an *approval-based committee (ABC) election*. An *ABC voting rule* takes as input an ABC election $(A, k)$ and outputs one or more committees of size $k$.

As is standard in the ABC election literature, we assume that voters only care about the number of approved candidates in the committee, i.e., voter $i$ evaluates a committee $W$ by $|A_i \cap W|$. Given a subset $S \subseteq C$ of candidates, we let $N(S)$ denote the set of voters who approve all candidates in $S$, i.e., $N(S) = \{i \in N \colon S \subseteq A_i\}$.

Given an ABC election $(A, k)$ and $\ell \in \mathbb{N}$, we call a group $V \subseteq N$ of voters *$\ell$-cohesive* if

$$|V| \geqslant \ell \cdot \frac{n}{k} \quad \text{and} \quad |\bigcap_{i \in V} A_i| \geqslant \ell.$$

The following representation notions are due to Aziz et al. [2017]. A committee $W \subseteq C$ of size $k$ provides

- *justified representation (JR)* if for each 1-cohesive group $V \subseteq N$, there is a voter $i \in V$ with $|W \cap A_i| \geqslant 1$;
- *extended justified representation (EJR)* if for each $\ell \in \mathbb{N}$ and each $\ell$-cohesive group $V \subseteq N$, there is a voter $i \in V$ with $|W \cap A_i| \geqslant \ell$; and
- *core stability* if for each group $V \subseteq N$ (independent of $V$ being $\ell$-cohesive) and $S \subseteq C$ with $|S| \leqslant |V| \cdot \frac{n}{k}$ there is a voter $i \in V$ with $|W \cap A_i| \geqslant |S \cap A_i|$.

All of these notions have in common that they consider a group of voters "represented" as long as at least one voter in the group is sufficiently represented. This point of view might be hard to justify in many contexts. In the following section, we present our approach to representation that takes into account every voter in a group individually.

## 5.3 Individual Representation (IR)

In this section, we define the main concept of this chapter: individual representation. This notion builds on the idea of (semi-)strong justified representation defined by Aziz et al. [2017] and the notion of individual fairness in clustering as defined by Jung et al. [2020].

Similarly to the proportionality notions defined above, we assume that a voter $i \in N$ deserves some representation in an ABC election if $i$ can find enough other voters who all approve a subset of candidates in common. This follows the rationale that everyone in a group of voters that makes up a sizable part of the electorate and can come to an agreement on how (part of) the committee ought to be filled, should be represented accordingly.

Given an ABC election $(A, k)$, we determine the number of seats that voter $i \in N$ can justifiably demand as

$$f_i \coloneqq \max_{S \subseteq A_i} \{|S| \colon |N(S)| \geqslant |S| \cdot n/k\}.$$

In words, $f_i$ is the largest value $f$ such that voter $i$ can find enough like-minded voters to form an $f$-cohesive group.

**Definition 8** (Individual Representation)**.** *Given an ABC election $(A, k)$, a committee $W \subseteq C$ of size $k$ provides* individual representation (IR) *if $|W \cap A_i| \geqslant f_i$ for all voters $i \in N$.*

When only requiring $|W \cap A_i| \geqslant 1$ for every voter with $f_i \geqslant 1$ we get *semi-strong justified representation (semi-strong JR)* as defined by Aziz et al. [2017]. The authors of that paper provide an example showing that semi-strong JR committees do not always exist. Since individual representation clearly is a more demanding property, it immediately follows that *IR committees* (i.e., committees providing IR) do not need to exist either.

IR not only implies semi-strong JR, but also EJR. More precisely, every IR committee also provides EJR. To build intuition on how our notion differs from established ones, and on how it leads to the election of committees that are "fair" from an individual voter perspective, consider the following examples, illustrated in Figure 5.1.

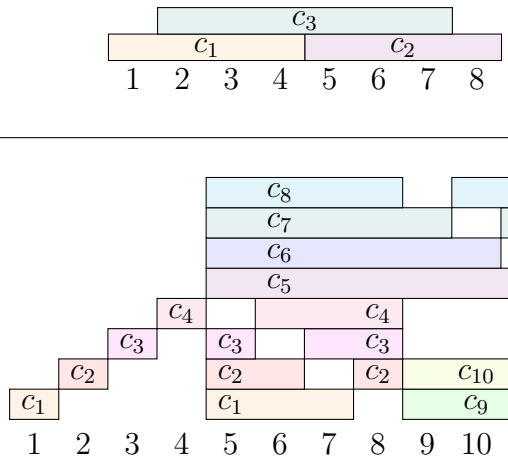


Figure 5.1: Two profiles admitting IR committees that are not identified by common voting rules or other representation axioms. Voters correspond to integers and approve all candidates placed above them.

**Example 2.** *The first part of Figure 5.1 shows an approval profile with* $8$ *voters and* $3$ *candidates. Assuming* $k = 2$*, every voter* $i \in N$ *has* $f_i = 1$*. Thus, the only committee providing IR is* $W = \{c_1, c_2\}$*, which represents every voter once and, moreover, satisfies core stability. However, both* $W' = \{c_1, c_3\}$ *and* $W'' = \{c_2, c_3\}$ *are core stable as well (and in fact would be selected when choosing a committee maximizing social welfare). Many common ABC voting rules would select either* $W'$ *or* $W''$ *(see Proposition 5). One can argue that committee* $W$ *is a much "fairer" or "more representative" choice in this example.*

**Example 3.** *The second part of Figure 5.1 shows an approval profile with* $12$ *voters and* $10$ *candidates. For* $k = 6$ *we compute* $f_i = 1$ *for* $i \in \{1, \ldots, 4\}$ *and* $f_i = 2$ *for* $i \in \{5, \ldots, 12\}$*. Here, the only committee providing IR is* $W = \{c_1, c_2, c_3, c_4, c_9, c_{10}\}$ *representing each of the first four voters once, while representing all other voters at least twice. This committee is not core stable, because the voter group* $\{5, \ldots, 12\}$ *would prefer* $\{c_5, c_6, c_7, c_8\}$ *to* $W$*. In order to appreciate the IR committee* $W$*, consider voters* $1$ *to* $4$ *and observe that these voters are completely "symmetric." Hence, from an "equal treatment of equals" perspective, if one of them is represented by an approved candidate in the committee, the same should hold for the others. In fact, the only core-stable committees that provide this kind of symmetry are* $\{c_5, \ldots, c_{10}\}$*, in which one third of the electorate is not represented at all, or committees containing only two candidates among* $c_5$ *to* $c_8$*. In the latter case, by noticing that voters* $9$ *to* $12$ *are "symmetric" as well, we can argue in a similar way as above that they are not treated equally. Thus, the committee* $W$ *that uniquely provides IR can be considered the "fairest" choice under an individualistic point of view.*

The instance in Example 3 shows that core stability and individual representation are incompatible, in the strong sense that the (nonempty) set of IR committees and the (nonempty) set of core-stable committees are disjoint.

**Proposition 3.** *IR is incompatible with core stability.*

Next, we show a further incompatibility result.

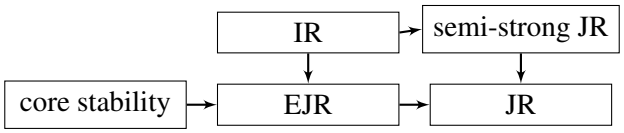


Figure 5.2: Relationships between individual representation (IR) and the representation notions introduced by Aziz et al. [2017]. An arrow from $X$ to $Y$ depicts that $X$ implies $Y$ (if $X$ is possible in the given instance).

**Proposition 4.** *Semi-strong JR is incompatible with EJR.*

*Proof.* Consider an ABC election with $n = 8$, $k = 4$ and the following approval profile: $A_1 = \cdots = A_4 = \{c_1, c_2\}$, $A_5 = \{c_3, c_4, c_5\}$, $A_6 = \{c_3\}$, $A_7 = \{c_4\}$, and $A_8 = \{c_5\}$. As $f_6 = f_7 = f_8 = 1$, every committee $W$ that provides semi-strong JR must satisfy $\{c_3, c_4, c_5\} \subseteq W$. But then we have $|W \cap A_i| < 2$ for all $i \in [4]$, even though these four voters form a 2-cohesive group. □

**Approximate notions.** As we mentioned above, IR is not always attainable. The immediate next question is whether we can guarantee IR in an approximate sense. To study this question, we introduce the notion of $(\alpha, \beta)$-individual representation, which uses additive *and* multiplicative approximation parameters.

**Definition 9** ($(\alpha, \beta)$-Individual Representation)**.** *Given an ABC election $(A, k)$, a committee $W$ provides $(\alpha, \beta)$-*individual representation ($(\alpha, \beta)$-IR) *if for every voter $i \in N$, $\alpha \cdot |A_i \cap W| + \beta \geqslant f_i$, with $\alpha \geqslant 1$ and $\beta \geqslant 0$.*

Unfortunately, non-trivial approximation guarantees are impossible to obtain without restricting the set of profiles.

**Theorem 18.** *There exists an instance that does not admit an $(\alpha, \beta)$-IR committee for $\beta < k - 1$, and any $\alpha \geqslant 1$.*

*Proof.* Assume that $n = k \cdot (k + 1)$. Then it holds that $n/k = k + 1 > k$. Consider a profile in which for each voter $i \in [n/k]$, $A_i = \{c_{(k-1)\cdot(i-1)+1}, \ldots, c_{(k-1)\cdot i}\}$ and all remaining $n - n/k$ voters approve all candidates. Thus, for every voter $i \in [n/k]$ we get that $|N(A_i)| = 1 + n - n/k = 1 + (k-1)(k+1)$. Since $n/k = k + 1$, this implies that $f_i \geqslant k - 1$. Further, for all distinct voters $i, i' \in [n/k]$ it holds that $|A_i \cap A_{i'}| = 0$. However, since $n/k > k$, for each $W \subseteq C$ with $|W| = k$ there is a voter $i \in [n/k]$ with $|A_i \cap W| = 0$. Thus, for any $\alpha \geqslant 1$ and $\beta < k - 1$ this instance does not admit a $(\alpha, \beta)$-IR committee. □

To see that this bound is worst-possible, note that if $f_i = k$ for some voter $i$, this means that *all* voters have a set of at least $k$ jointly approved candidates (and a committee consisting of such candidates would provide IR). On the other hand, every committee trivially provides $(1, k - 1)$-IR whenever $f_i < k$ for all $i \in N$.

### 5.3.1 ABC Rules Violating IR

While the inapproximability result in the previous section is quite discouraging, a natural follow-up question is whether we can at least find IR committees *whenever they exist*. This question was already raised by Aziz et al. [2017] in the context of semi-strong JR, but remained open. In other words, we look for rules that are "consistent" with individual representation.

**Definition 10** (IR-consistency)**.** *An ABC rule is* consistent with individual representation*, or short* IR-consistent*, if it outputs at least one IR committee for every ABC election that admits one.*

Consistency with semi-strong JR can be defined analogously. We show that all common ABC voting rules fail consistency with respect to both IR and semi-strong JR.[2]

Example 2 already rules out any rule that always selects one of the candidates with the highest numbers of approvals, so-called *approval winners*. In particular, this class of rules includes all commmon committee-monotonic ABC rules as well as other "sequential" rules like Rule X, as these rules select one of the approval winners in the very first round.

**Proposition 5.** *No ABC voting rule that always selects one of the approval winners into the winning committee is IR-consistent.*

Moreover, the rules PAV, Satisfaction-AV, and reverse-seqPAV select only committees including $c_3$ in Example 2, and thus fail IR-consistency. Below, we provide additional examples showing that all remaining ABC rules mentioned in Table 1 of the survey by Lackner and Skowron [2021] fail IR-consistency as well.

**Example 4.** *Consider the following profile with $n = 16$ voters and assume $k = 4$:*

$$\begin{aligned} A_1 = A_2 = A_3 &= \{c_1\} & A_4 &= \{c_1, c_5\} \\ A_5 = A_6 = A_7 &= \{c_2\} & A_8 &= \{c_2, c_5\} \\ A_9 = A_{10} = A_{11} &= \{c_3\} & A_{12} &= \{c_3, c_5\} \\ A_{13} = A_{14} = A_{15} &= \{c_4\} & A_{16} &= \{c_5\}. \end{aligned}$$

*Here we have $f_i = 1$ for all voters $i \in N \setminus \{13, 14, 15\}$ and thus the only committee providing individual representation is $W = \{c_1, c_2, c_3, c_5\}$. Chamberlin-Courant-AV (CCAV), Monroe-AV and PAV with the weight-vector $(1, \frac{1}{n}, \frac{1}{n^2}, \ldots)$, which provide an IR committee in the example of Proposition 5, choose $\{c_1, c_2, c_3, c_4\}$ and thus fail individual representation.*

The only two remaining rules from Table 4.1 in the survey by Lackner and Skowron [2023] are leximin-Phragmén [Brill et al., 2017] and Minimax-AV [Brams et al., 2007].

**Example 5.** *Consider the following profile with $n = 6$ voters and let $k = 3$:*

$$\begin{aligned} A_1 &= \{c_1\} & A_3 &= \{c_1, \ldots, c_5\} & A_5 &= \{c_4, c_5, c_6\} \\ A_2 &= \{c_2\} & A_4 &= \{c_3, \ldots, c_6\} & A_6 &= \{c_5, c_6\}. \end{aligned}$$

[2]Since neither IR nor semi-strong JR is always achievable (and semi-strong JR may be achievable in instances where IR is not) we can, in general, not deduce consistency regarding one of the notions from consistency regarding the other. However, all our examples satisfy $f_i \leqslant 1$ for all voters $i$, such that semi-strong JR and IR coincide.

*Here we have $f_i = 1$ for all voters $i \in N$ and thus $\{c_1, c_2, c_5\}$ and $\{c_1, c_2, c_6\}$ are the only committees providing individual representation (or semi-strong JR). It turns out that leximin-Phragmén does not select any of these two. It is easy to see that the load distribution with the minimum maximal load for either of the two IR committees is $(\frac{2}{3}, \frac{2}{3}, \frac{2}{3}, \frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. The committee $\{c_1, c_4, c_5\}$, however, induces a load distribution of $(\frac{3}{5}, 0, \frac{3}{5}, \frac{3}{5}, \frac{3}{5}, \frac{3}{5})$ which is lower both in terms of the maximum as well as lexicographic ordering.*

**Example 6.** *Consider the following profile with $n = 100$ voters and let $k = 2$:*

$$99 \times \{c_1, c_2\} \qquad 1 \times \{c_3, \ldots, c_8\}.$$

*Minimax-AV (which minimises the maximum Hamming-distance among all voters to the winning committee) selects any two candidates from $\{c_3, \ldots c_8\}$ and none that is supported by the 99 voters. This clearly violates individual representation and semi-strong JR.*

### 5.3.2 Computational Complexity

Another open problem stated by Aziz et al. [2017] concerns the computational complexity of deciding whether a given ABC election admits a committee providing semi-strong JR. We settle this question and the analogous one for individual representation by showing that both problems are NP-hard.

**Theorem 19.** *It is NP-hard to decide whether an ABC election admits an IR committee or a semi-strong JR committee.*

*Proof.* We reduce from *exact cover by 3-sets*. Here, we are given a set of elements $X = \{x_1, \ldots, x_{3\ell}\}$ and a set of sets $\mathcal{T} \subseteq 2^X$ such that $|T| = 3$ for all $T \in \mathcal{T}$. The goal is to find a partition of $X$ into sets from $\mathcal{T}$. The problem is NP-hard even if each element appears in exactly three sets [Garey and Johnson, 1979]. We construct an ABC instance by setting $N = X$ and $C = \mathcal{T}$. Further, each set $T \in \mathcal{T} = \{x_i, x_j, x_l\}$ is approved exactly by voters $x_i, x_j$, and $x_l$. We set $k = \ell$. Hence, only groups of 3 voters corresponding to sets in $\mathcal{T}$ are 1-cohesive, and we get $f_{x_i} = 1$ for each $x_i \in X$.

Every exact cover by 3-sets corresponds to a committee of size $k$ where every voter is represented exactly once and thus provides IR in this instance. Conversely, every IR committee of the constructed ABC instance corresponds to a selection of sets from $\mathcal{T}$ such that every element in $X$ is covered exactly once. Since $f_i = 1$ for every voter, the same argument holds for semi-strong JR as well. □

Moreover, it is hard to decide whether a given voter's $f_i$-value exceeds a given value.

**Theorem 20.** *Given an ABC instance, a voter $i$, and $j \in \mathbb{N}$, it is NP-complete to decide whether $f_i \geqslant j$.*

*Proof.* It is easy to see that this problem is in NP since any subset of voters including voter $i$ of size $j \cdot \frac{|N|}{k}$ and any subset of candidates of size $j$ approved by all selected voters serves as a witness.

We reduce from *balanced complete bipartite subgraph*. Here, we are given a bipartite graph $G = (V_1 \cup V_2, E)$ and an integer $j$ and the goal is to decide whether $G$ has $K_{j,j}$ as a subgraph, i.e., a subgraph consisting of $j$ vertices from $V_1$ and $j$ vertices from $V_2$ forming a bipartite clique. The problem is known to be NP-hard [Garey and Johnson, 1979]. We construct an ABC instance by setting $N = V_1 \cup \{x\}$, $C = V_2 \cup \{y\}$ and $k = |V_1| + 1$. Thus, $\frac{n}{k} = 1$. Each $v \in V_1$ approves exactly its

neighbors in $G$, as well as $y$, while $x$ approves all candidates. It follows that $f_x \geqslant j+1$ if and only if there is a set of $j$ voters different from $x$ approving at least a common set of $j+1$ candidates. Since all voters approve $y$, this is equivalent to these $j$ voters all approving $j$ candidates different from $y$ and therefore by definition all being connected to these $j$ vertices in $V_2$. Thus, they form a $K_{j,j}$ if and only if $f_x \geqslant j+1$. □

## 5.4 Domain Restrictions

We have seen in Theorem 18 that non-trivial approximations of individual representation are impossible to obtain in general. In this section, we explore whether this negative result can be circumvented by considering restricted domains of preferences. Domain restrictions for dichotomous (i.e., approval) preferences have been studied by Elkind and Lackner [2015] and Yang [2019].

Restricting attention to a well-structured domain often allows for axiomatic and algorithmic results that are not achievable otherwise [Elkind et al., 2017]. In the ABC setting, for example, it has recently been shown that a core-stable committee always exists in certain restricted domains [Pierczyński and Skowron, 2022], whereas the existence of such committees is an open problem for the unrestricted domain.

We start by recalling the definitions of two classic restricted domains of dichotomous preferences: *candidate interval* and *voter interval* [Elkind and Lackner, 2015].

**Definition 11** (Candidate Interval)**.** *An approval profile $A$ satisfies* candidate interval (CI) *if there is a linear order over the candidates $C$ such that for every voter $i \in N$, $A_i$ consists of an interval of that order.*

**Definition 12** (Voter Interval)**.** *An approval profile $A$ satisfies* voter interval (VI) *if there is a linear order over the voters $N$ such that for every candidate $c_j \in C$, the voters that approve $c_j$ build an interval of that order.*

The profile in Example 2 satisfies *both* candidate interval and voter interval. In fact, a voter order witnessing VI is given in Figure 5.1. To see that the profile satisfies CI as well, consider the order $(c_1, c_3, c_2)$. The profile in Example 3, on the other hand, satisfies neither CI nor VI.

Elkind and Lackner [2015] have shown that it can be checked in polynomial time whether a profile satisfies CI or VI. (If the answer is yes, a linear order over candidates/voters can be found efficiently as well.)

Our first observation is that the *candidate interval* domain is not helpful for our purposes: Indeed, the approval profile used to establish Theorem 18 can easily be seen to satisfy CI. Thus, restricting preferences in this way does not yield any improved bounds.

**Corollary 21.** *There exists a CI profile that does not admit an $(\alpha, \beta)$-IR committee with $\beta < k-1$ and any $\alpha \geqslant 1$.*

Now, we turn our attention to the *voter interval* domain. Due to the similarity between VI and CI, one might expect a similar result here. Surprisingly, however, we can prove a positive result for VI: We provide an algorithm that finds a $(2, 4)$-IR committee in polynomial time for any VI profile.

**Algorithm 4** $(2, 4)$-IR for Voter Interval Profiles

$W_0 \leftarrow \emptyset$ {Round 1}
**for** $i = 1$ to $n$ **do**
  **if** $|W_{i-1} \cap A_i| < \lfloor |N_{\geqslant i}| \cdot k/(2n) \rfloor$ **then**
    Let $S_i$ be an arbitrary subset of $S_i^*$ of size equal to $\lfloor |N_{\geqslant i}| \cdot k/(2n) \rfloor - |W_{i-1} \cap A_i|$ such that $|S_i \cap W_{i-1}|$ is minimized
    $W_i \leftarrow W_{i-1} \cup S_i$
  **end if**
**end for**
$\widehat{W}_1 \leftarrow \emptyset$ {Round 2}
**for** $i = n$ to $1$ **do**
  **if** $|\widehat{W}_{n-i} \cap A_i| < \lfloor |N_{<i}| \cdot k/(2n) \rfloor$ **then**
    Let $S_i$ be an arbitrary subset of $S_i^*$ of size equal to $\lfloor |N_{<i}| \cdot k/(2n) \rfloor - |\widehat{W}_{n-i} \cap A_i|$ such that $|S_i \cap (\widehat{W}_{n-i} \cup W_n)|$ is minimized
    $\widehat{W}_{n-i+1} \leftarrow \widehat{W}_{n-i} \cup S_i$
  **end if**
**end for**
$S \leftarrow$ an arbitrary subset of $C$ with $S \cap (W_n \cup \widehat{W}_n) = \emptyset$ and $|S| = k - |W_n| - |\widehat{W}_n|$
**return** $W_n \cup \widehat{W}_n \cup S$

Before describing the high level idea of our algorithm we state a useful property of VI profiles. Without loss of generality, we assume that the linear order witnessing VI is given by $(1, \dots, n)$. Moreover, for $a, b \in \mathbb{Z}$ with $a \leqslant b$, we let $[a, b]$ denote the integer interval $\{a, a + 1, \dots, b\}$.

**Observation 22.** *Let $i_1, i_2, i_3 \in [n]$ such that $i_1 < i_2 < i_3$. For any $S \subseteq C$, if $i_1 \in N(S)$ and $i_3 \in N(S)$, then $i_2 \in N(S)$.*

Let $S_i^* \coloneqq \arg\max_{S \subseteq A_i} \{|S| \colon |N(S)| \geqslant |S| \cdot n/k\}$, i.e., a largest subset of $A_i$ approved by sufficiently many voters to validate the $f_i$-value. (If multiple such sets exist, ties are broken arbitrarily.) From Observation 22 we know that if $i_1, i_2, i \in [n]$ such that $i_1 < i_2 < i$ or $i < i_2 < i_1$, and if $i_1 \in N(S_i^*)$, then $i_2 \in N(S_i^*)$, i.e., $N(S_i^*)$ forms an interval of the order of voters including $i$.

Further, let $N_{<i} \coloneqq \{i' \in N(S_i^*) \colon i' \leqslant i - 1\}$ denote the set of voters in $N(S_i^*)$ that are ordered *before* $i$ and let $N_{\geqslant i} \coloneqq \{i' \in N(S_i^*) \colon i' \geqslant i\}$ denote the set of voters in $N(S_i^*)$ that are ordered *after* $i$ (including $i$ itself).

**Observation 23.** *For each voter $i$, there exist $a_i \in [1, i]$ and $b_i \in [i, n]$ such that $N_{<i} = [a_i, i - 1]$ and $N_{\geqslant i} = [i, b_i]$.*

Using this observation and the fact that $f_i = |S_i^*| \leqslant (|N_{<i}| + |N_{\geqslant i}|) \cdot k/n$, Algorithm 4 returns a $(2, 4)$-IR committee $W$ for any VI profile as follows. In the first round, iterating from voter $i = 1$ to $n$, it selects at least $\lfloor |N_{\geqslant i}| \cdot k/(2n) \rfloor$ candidates that are approved by voter $i$. In the second round, iterating from voter $i = n$ to 1, it selects at least $\lfloor |N_{<i}| \cdot k/(2n) \rfloor$ candidates that are approved by voter $i$ (excluding the candidates that are selected in the first round). Together, this ensures $|W \cap A_i| \geqslant f_i/2 - 2$, where $W$ is the set of selected candidates.

**Theorem 24.** *For every approval profile satisfying voter interval, Algorithm 4 returns a* $(2,4)$*-IR committee in polynomial time.*

*Proof.* Let $W = W_n \cup \widehat{W}_n \cup S$ be the committee returned by Algorithm 4, and let $f_{\geqslant i} = |N_{\geqslant i}| \cdot k/n$ and $f_{<i} = |N_{<i}| \cdot k/n$. In the first round we ensure that $|W_n \cap A_i| \geqslant \lfloor f_{\geqslant i}/2 \rfloor$, as at iteration $i$ if $|W_{i-1} \cap A_i| < \lfloor f_{\geqslant i}/2 \rfloor$, we include $\lfloor f_{\geqslant i}/2 \rfloor - |W_{i-1} \cap A_i|$ candidates into $W_i$ that are not already included. Similarly, in the second round we ensure that $|\widehat{W}_n \cap A_i| \geqslant \lfloor f_{<i}/2 \rfloor$ as at iteration $n-i+1$ if $|\widehat{W}_{n-i} \cap A_i| < \lfloor f_{\geqslant i}/2 \rfloor$, we include $\lfloor f_{\geqslant i}/2 \rfloor - |\widehat{W}_{n-i} \cap A_i|$ candidates that are not already included. As $f_i \leqslant f_{\geqslant i} + f_{<i}$, for each $i \in N$ we have that

$$|W \cap A_i| \geqslant \lfloor f_{\geqslant i}/2 \rfloor + \lfloor f_{<i}/2 \rfloor \geqslant f_i/2 - 2,$$

and therefore $2 \cdot |W \cap A_i| + 4 \geqslant f_i$. Thus, we conclude that $W$ provides $(2,4)$-IR.

Now we show that $|W_n| \leqslant k/2$ and $|\widehat{W}_n| < k/2$. We first consider $W_n$.

**Lemma 16.** $|W_i| \leqslant \frac{((i-1)+|N_{\geqslant i}|)\cdot k}{2n}$ *for all* $i \in [n]$.

*Proof.* We prove the lemma using induction. For $i = 1$, $W_1 = \lfloor f_{\geqslant 1}/2 \rfloor \leqslant \frac{|N_{\geqslant 1}|\cdot k}{2n}$ and the statement holds. Assume that for all $t' < t$, we have $|W_{t'}| \leqslant \frac{((t'-1)+|N_{\geqslant t'}|)\cdot k}{2n}$.

We show that the statement holds for $W_t$. Note that

$$|W_t| = |W_{t-1}| + \lfloor f_{\geqslant t}/2 \rfloor - |W_{t-1} \cap A_t| \tag{5.1}$$

as at iteration $t$, the algorithm adds $\lfloor f_{\geqslant t}/2 \rfloor - |W_{t-1} \cap A_t|$ candidates to $W_t$. Let $t^* = \max\{r \in [t-1] \colon t \in N(S^*_{t-r})\}$. If $t^* = 0$, from Equation (5.1) we get that

$$\begin{aligned}
|W_t| \leqslant & \frac{((t-2)+|N_{\geqslant t-1}|)\cdot k}{2n} \\
& + \lfloor f_{\geqslant t}/2 \rfloor - |W_{t-1} \cap A_t| \\
\leqslant & \frac{((t-2)+1)\cdot k}{2n} + \lfloor f_{\geqslant t}/2 \rfloor - |W_{t-1} \cap A_t| \\
\leqslant & \frac{(t-1+|N_{\geqslant t}|)\cdot k}{2n},
\end{aligned}$$

where the last inequality follows from the fact that $|N_{\geqslant t-1}| = 1$, as $t \notin N(S^*_{t-1})$. Now assume that $t^* > 0$. First, using induction, we show that

$$|W_t| = |W_{t-r}| + \lfloor f_{\geqslant t}/2 \rfloor - |W_{t-r} \cap A_t|$$

for any $r \in [1, t^*+1]$.

For $r = 1$, the claim immediately follows from Equation (5.1). Assume that for all $q' < q$ it holds that $|W_t| = |W_{t-q'}| + \lfloor f_{\geqslant t}/2 \rfloor - |W_{t-q'} \cap A_t|$. Now, we have

$$\begin{aligned}
|W_{t-(q-1)} \cap A_t| = & \left\lfloor f_{\geqslant t-(q-1)}/2 \right\rfloor - |W_{t-q} \cap A_{t-(q-1)}| \\
& + |W_{t-q} \cap A_t|
\end{aligned} \tag{5.2}$$

as at iteration $t-(q-1)$, $\lfloor f_{\geqslant t-(q-1)}/2 \rfloor - |W_{t-q} \cap A_{t-(q-1)}|$ candidates from $S^*_{t-(q-1)}$ are added to $W_{t-(q-1)}$, and as $t \in N(S^*_{t-(q-1)})$, these candidates are approved by $t$, too. Then,

$$\begin{aligned}
|W_t| &= |W_{t-(q-1)}| + \lfloor f_{\geqslant t}/2 \rfloor - |W_{t-(q-1)} \cap A_t| \\
&= |W_{t-q}| + \lfloor f_{\geqslant t-(q-1)}/2 \rfloor - |W_{t-q} \cap A_{t-(q-1)}| \\
&\quad + \lfloor f_{\geqslant t}/2 \rfloor - |W_{t-(q-1)} \cap A_t| \\
&= |W_{t-q}| + \lfloor f_{\geqslant t}/2 \rfloor - |W_{t-q} \cap A_t|,
\end{aligned}$$

where the second transition follows since at iteration $t-(q-1)$, $\lfloor f_{\geqslant t-(q-1)}/2 \rfloor - |W_{t-q} \cap A_{t-(q-1)}|$ candidates are added to $W_{t-(q-1)}$, and the third transition follows from Equation (5.2).

Now, we distinguish two cases.

**Case 1:** $t^* = t-1$**.** Here, we have

$$\begin{aligned}
|W_t| &= |W_{t-(t^*+1)}| + \lfloor f_{\geqslant t/2} \rfloor - |W_{t-(t^*+1)} \cap A_t| \\
&= |W_0| + \lfloor f_{\geqslant t/2} \rfloor - |W_0 \cap A_t| = \frac{|N_{\geqslant t}| \cdot k}{2n}.
\end{aligned}$$

**Case 2:** $t^* > 1$**.** Here, we have

$$\begin{aligned}
|W_t| =& |W_{t-(t^*+1)}| + \lfloor f_{\geqslant t/2} \rfloor - |W_{t-(t^*+1)} \cap A_t| \\
\leqslant& \frac{(t-(t^*+1)-1+|N_{\geqslant t-(t^*+1)}|) \cdot k}{2n} \\
&+ \frac{|N_{\geqslant t}| \cdot k}{2n} \\
\leqslant& \frac{(t-1+|N_{\geqslant t}|) \cdot k}{2n},
\end{aligned}$$

where the third inequality follows from the fact that $|N_{\geqslant t-(t^*+1)}| \leqslant t-(t-(t^*+1))$, as $t$ is not in $N(S^*_{t-(t^*+1)})$. □

As Rounds 1 and 2 of Algorithm 4 are symmetric, with similar arguments we can show the following lemma.

**Lemma 17.** $|\widehat{W}_i| \leqslant \frac{((i-1)+|N_{<n-i+1}|) \cdot k}{2n}$ *for all* $i \in [n]$.

*Proof.* We prove the lemma using induction. For $i=1$, we have that $\widehat{W}_1 = \lfloor f_{<n}/2 \rfloor \leqslant \frac{|N_{<n}| \cdot k}{2n}$ and the statement holds. Assume that, for all $t' < t$, we have $|\widehat{W}_{t'}| \leqslant \frac{((t'-1)+|N_{<n-t'+1}|) \cdot k}{2n}$.

We show that the statement holds for $\widehat{W}_t$. Note that

$$|\widehat{W}_t| = |\widehat{W}_{t-1}| + \lfloor f_{<n-t+1}/2 \rfloor - |\widehat{W}_{t-1} \cap A_{n-t+1}| \tag{5.3}$$

as at iteration $t$, the algorithm adds $\lfloor f_{<n-t+1}/2 \rfloor - |\widehat{W}_{t-1} \cap A_{n-t+1}|$ candidates to $\widehat{W}_t$. Let $t^* = \max\{r \in [t-1] \colon n-t+1 \in N(S^*_{n-t+1+r})\}$. If $t^* = 0$, from Equation (5.3) we get that

$$\begin{aligned}
|\widehat{W}_t| &\leqslant \frac{((t-2)+|N_{<n-t+2}|)\cdot k}{2n} + \lfloor f_{<n-t+1}/2 \rfloor - |\widehat{W}_{t-1} \cap A_{n-t+1}| \\
&\leqslant \frac{((t-2)+0)\cdot k}{2n} + \lfloor f_{<n-t+1}/2 \rfloor - |\widehat{W}_{t-1} \cap A_{n-t+1}| \\
&\leqslant \frac{((t-2)+|N_{<n-t+1}|)\cdot k}{2n},
\end{aligned}$$

where the second inequality follows from the fact that $|N_{<n-t+2}| = 0$, as $n-t+1 \notin N(S^*_{n-t+2})$. Now assume that $t^* > 0$. First, using induction, we show that

$$|\widehat{W}_t| = |\widehat{W}_{t-r}| + \lfloor f_{<n-t+1}/2 \rfloor - |\widehat{W}_{t-r} \cap A_{n-t+1}|$$

for any $r \in [1, t^*+1]$.

For $r = 1$, the claim immediately follows from Equation (5.3). Assume that for all $q' < q$ it holds that $|\widehat{W}_t| = |\widehat{W}_{t-q'}| + \lfloor f_{<n-t+1}/2 \rfloor - |\widehat{W}_{t-q'} \cap A_{n-t+1}|$. Now, we have

$$\begin{aligned}
|\widehat{W}_{t-(q-1)} \cap A_{n-t+1}| = \big\lfloor f_{<n-t+(q-1)+1}/2 \big\rfloor &- |\widehat{W}_{t-q} \cap A_{n-t+(q-1)+1}| \\
&+ |\widehat{W}_{t-q} \cap A_{n-t+1}| \qquad (5.4)
\end{aligned}$$

as at iteration $t-(q-1)$ we add $\big\lfloor f_{<n-t+(q-1)+1}/2 \big\rfloor - |\widehat{W}_{t-q} \cap A_{n-t+(q-1)+1}|$ candidates from $S^*_{n-t+(q-1)+1}$ to $\widehat{W}_{t-(q-1)}$, and as $n-t+1 \in N(S^*_{n-t+(q-1)+1})$, these candidates are approved by $n-t+1$, too. Then,

$$\begin{aligned}
|\widehat{W}_t| &= |\widehat{W}_{t-(q-1)}| + \lfloor f_{<n-t+1}/2 \rfloor - |\widehat{W}_{t-(q-1)} \cap A_{n-t+1}| \\
&= |\widehat{W}_{t-q}| + \big\lfloor f_{<n-t+(q-1)+1}/2 \big\rfloor - |\widehat{W}_{t-q} \cap A_{n-t+(q-1)+1}| \\
&\qquad + \lfloor f_{<n-t+1}/2 \rfloor - |\widehat{W}_{t-(q-1)} \cap A_{n-t+1}| \\
&= |\widehat{W}_{t-q}| + \lfloor f_{<n-t+1}/2 \rfloor - |\widehat{W}_{t-q} \cap A_{n-t+1}|,
\end{aligned}$$

where the second transition follows since at iteration $t-(q-1)$, $\big\lfloor f_{<n-t+(q-1)+1}/2 \big\rfloor - |\widehat{W}_{t-q} \cap A_{n-t+(q-1)+1}|$ candidates are added to $\widehat{W}_{t-(q-1)}$, and the third transition follows from Equation (5.4).

Now, we distinguish two cases.

**Case 1:** $t^* = t-1$ Here, we have

$$\begin{aligned}
|\widehat{W}_t| &= |\widehat{W}_{t-(t^*+1)}| + \big\lfloor f_{<n-t+1/2} \big\rfloor - |\widehat{W}_{t-(t^*+1)} \cap A_{n-t+1}| \\
&= |\widehat{W}_0| + \lfloor f_{<n-t+1}/2 \rfloor - |\widehat{W}_0 \cap A_{n-t+1}| = \frac{|N_{<n-t+1}|\cdot k}{2n}.
\end{aligned}$$

**Case 2:** $t^* > 1$ Here, we have

$$\begin{aligned}
|\widehat{W}_t| &= |\widehat{W}_{t-(t^*+1)}| + \lfloor f_{<n-t+1}/2 \rfloor - |\widehat{W}_{t-(t^*+1)} \cap A_{n-t+1}| \\
&\leqslant \frac{(t-(t^*+1)-1+|N_{<n-t+(t^*+1)+1}|)\cdot k}{2n} + \frac{|N_{<n-t+1}|\cdot k}{2n} \\
&\leqslant \frac{(t-2+|N_{<n-t+1}|)\cdot k}{2n},
\end{aligned}$$

**Algorithm 5** Finding $f_i$ and $S_i^*$

$f_i \leftarrow 0$
$S_i^* \leftarrow \emptyset$
**for** $a_i = 1$ to $i$ **do**
  **for** $b_i = i$ to $n$ **do**
    $S \leftarrow \emptyset$
    **for** $j = 1$ to $m$ **do**
      **if** $\{a_i, \ldots, b_i\} \subseteq N(\{c_j\})$ **then**
        $S \leftarrow S \cup \{c_j\}$
      **end if**
    **end for**
    $\ell^* \leftarrow \arg\max_{\ell \in \mathbb{N}} \{\ell : b_i - a_i + 1 \geqslant \ell \cdot n/k\}$
    **if** $\min\{\ell^*, |S|\} > f_i$ **then**
      $S_i^* \leftarrow$ an arbitrary subset of $S$ of size $\min\{\ell^*, |S|\}$
      $f_i \leftarrow |S_i^*|$
    **end if**
  **end for**
**end for**
**return** $f_i$, $S_i^*$

where the third inequality follows from the fact that $|N_{<n-t+(t^*+1)+1}| \leqslant t^*$, as $n - t + 1$ is not in $N(S^*_{n-t+(t^*+1)+1})$. □

From Lemma 16, for $i = n$, we have $|N_{\geqslant n}| \leqslant 1$, and hence

$$|W_n| \leqslant \frac{((n-1)+1) \cdot k}{2n} \leqslant k/2.$$

From Lemma 17, for $i = n$, we have that $|N_{<1}| = 0$, and hence

$$|W_n| \leqslant \frac{((n-1)+0) \cdot k}{2n} < k/2.$$

Thus, $|W| = |W_n| + |\widehat{W_n}| < k$.

Lastly, we show that $S_i^*$, and thus $f_i$, can be computed in polynomial time. For this we employ Observation 23, i.e., the fact that $N(S_i^*)$ forms an interval of voters that includes $i$.

**Lemma 18.** *For any voter $i$, $f_i$ and $S_i^*$ can be computed in polynomial time.*

*Proof.* We show that Algorithm 5 correctly computes both $f_i$ and $S_i^*$. For each interval $[a_i, b_i]$ where $a_i \leqslant i$ and $i \leqslant b_i$, the algorithm finds the maximum number of candidates that this interval is eligible to elect, denoted by $\ell^*$. Moreover, $S \subseteq C$ denotes the set of candidates that is approved by all the voters in the interval. The algorithm calculates the maximum subset of $S$ that can be elected by the voters in the interval as $\min\{\ell^*, |S|\}$. Then, $S_i^*$ and $f_i$ are updated properly by assigning them the biggest subset and the size of it, respectively, that an interval of voters can elect.

Assume for contradiction that the algorithm returns a subset $S'$ with $|S'| < |S_i^*|$. This means that there is an interval $[a_i, b_i]$ of voters that can elect $S_i^*$. Thus, when the algorithm considers this interval, it would return a subset of size at least $|S_i^*|$, a contradiction. □

This concludes the proof of Theorem 24. □

Further, we can show that the bound provided by Theorem 24 is almost tight up to the additive part of 4.

**Theorem 25.** *There exists a VI approval profile that does not admit an $(\alpha, 0)$-IR committee with $\alpha < 2 - 2/k$.*

*Proof.* Consider the following instance with $k > 2$ and $m = 2(k-1)$. All the voters $i \in [2, n-1]$ approve all the candidates, while $A_1 = \{c_1, \dots, c_{k-1}\}$ and $A_n = \{c_{m-k+2}, \dots, c_m\}$. Notice that this profile is VI. Indeed, if we order the voters as $1, 2, \dots, n$, then the voters that approve each candidate form an interval of the ordering. Now, we can see that $f_1 = f_n = k-1$, but for each $W \subseteq C$, with $|W| = k$, either $|A_1 \cap W| \leqslant k/2$ or $|A_n \cap W| \leqslant k/2$. □

Next, we show that many common ABC rules are not guaranteed to return a committee that provides $(2, 4)$-IR for VI preferences.

**Example 7.** *Consider the following profile with $n = 28$ voters and let $k = 14$:*

$$\begin{aligned} A_1 &= \{c_1, \dots, c_7\} \\ A_2 = \dots = A_{14} &= \{c_1, \dots, c_{20}\} \\ A_{15} = \dots = A_{28} &= \{c_8, \dots, c_{20}\}. \end{aligned}$$

*Notice that $f_1 = 7$ and $f_i = 13$ for $i > 1$. Here, Rule X, AV, PAV, seq-PAV, rev-seq-PAV, seq-Phragmén, SAV, and Greedy Monroe (depending on the tie breaking) all return committees $W$ where $\{c_8, \dots, c_{20}\} \subset W$ and thus only one of the candidates among $\{c_1, \dots, c_7\}$ is included. Hence, they all fail to select a $(2, 4)$-IR committee.*

**Example 8.** *Consider the following profile with $n = 28$ voters and let $k = 14$:*

$$\begin{aligned} A_1 = \dots = A_{14} &= \{c_1, \dots, c_7\} \\ A_{10+i} &= \{c_{5+i}\}, \qquad \forall i \in \{1, \dots, 14\}. \end{aligned}$$

*Again, we have $f_1 = 7$. CC and seq-CC return committees $W$ where only one candidate among $\{c_1, \dots, c_7\}$ is included. Thus, they fail to select a $(2, 4)$-IR committee.*

A simple adaption of Example 6 shows that Minimax-AV fails to provide any good approximation for IR even on $t$-PART instances. From Table 4.1 of the recent survey by Lackner and Skowron [2023], the only rules missing are leximin-Phragmén and Monroe-AV for which we experimentally found profiles with $n = 30$ voters and $m = 300$ candidates where these two rules also fail to select a $(2, 4)$-IR committee. Thus, we can conclude that Algorithm 4 outperforms all common ABC voting rules in terms of approximating IR on VI profiles.

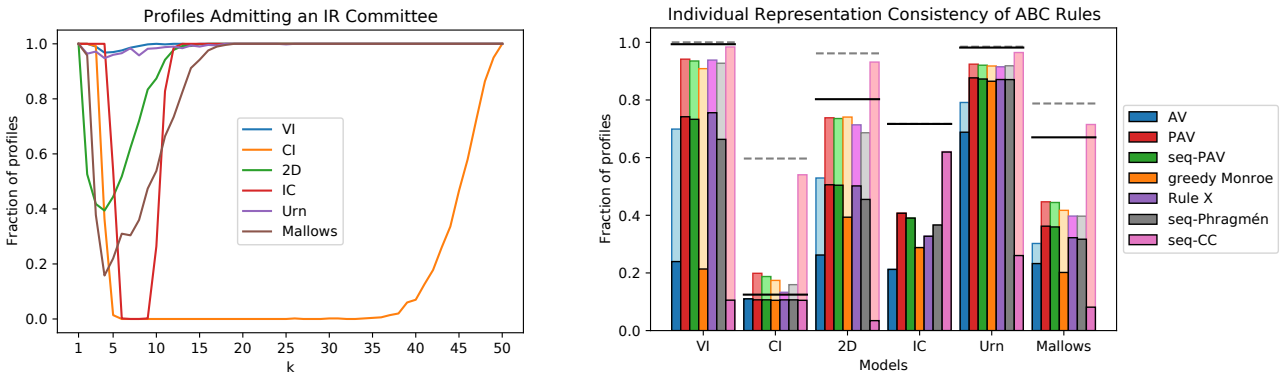


Figure 5.3: The figure on the left shows the ratio of generated profiles that admit an IR committee. In the figure on the right, for each model and each voting rule, the bold colored part of the bar indicates the ratio of instances the rule returned an IR committee, while the pale-colored part indicates the same ratio for semi-strong JR, averaged over all values $k$ with $2 \leqslant k \leqslant 20$. For each model, the black line indicates the fraction of instances admitting an IR committee, while the gray dashed line indicates the ratio of instances admitting a semi-strong JR committee.

## 5.5 Experimental Results

To complement our theoretical results, we performed experiments on generated approval profiles.

**Setup.** Inspired by Peters et al. [2021a], we used six models to generate approval profiles: a voter-interval model (VI), a candidate-interval model (CI), a two-dimensional euclidean model (2D), an impartial culture model (IC), an urn model (Urn), and a Mallows model (Mallows). We considered the ABC rules AV, PAV, seq-PAV, greedy Monroe, Rule X, seq-Phragmén, and sequential Chamberlin–Courant (seq-CC). For all conducted experiments, we have $100$ voters and $50$ candidates and generated $500$ instances for each model and for each $k \in [50]$.

**Results.** First, we studied how often the generated approval profiles admit an IR committee. The results are shown in the left graph of Figure 5.3. We found that IR committees exist quite often, especially for larger values of $k$. In particular, profiles generated by the VI model or by the urn model almost always admit IR committees. On the other hand, profiles generated by the CI model almost never admit IR committees (except for $k \geqslant 35$). This striking contrast between VI and CI, which is reminiscent of our theoretical results in Section 5.4, can be explained with a feature of the preference generation model: Due to the way we generate CI preferences, almost $20\%$ of voters have an approval set spanning at least $20\%$ of the candidates. These voters approving many candidates are then part of multiple cohesive groups, not all of which can be represented in an IR manner. (A similar situation can be observed in the profile that proves Theorem 18.)

Second, we studied how often different ABC rules select a committee providing IR (or semi-strong JR). In order not to dilute our results, we restricted $k$ to the "interesting" range between $2$ and $20$. The results are shown in the right graph of Figure 5.3. Of course, the fraction of profiles for which a rule selects an IR (or semi-strong JR) committee is upper-bounded by the fraction of profiles that admit such

a committee. For each model, the latter fraction is depicted in the graph as a solid black line for IR, and a dashed gray line for semi-strong JR. While no rule manages to find an IR committee every time one exists, the rules PAV, sequential PAV, sequential Phragmén, and Rule X select IR committees often. For the small fraction of CI profiles that admit an IR committee, all considered rules do a good job in finding one. Since seq-CC greedily optimizes the amount of voters that are represented at least once, it finds a committee providing semi-strong JR in almost all profiles that admit one. But as the rule does not aim at representing voters more than once, it almost never produces IR committees. In the profiles generated by the IC model, IR often coincides with semi-strong JR (for $k \leqslant 20$) because almost all non-zero $f_i$-values are $1$. This is in line with the effect noticed by Bredereck et al. [2019], whose experiments showed that EJR and JR are very likely to coincide under IC.

## 5.6 Discussion

Based on the observations that common axioms in approval-based committee voting do not address the representation of individual voters, and that common voting rules sometimes unfairly distinguish between such voters, we formalize *individual representation* (IR) as a requirement for committees. We find that all common voting rules fail to provide IR committees even when these exist. Nevertheless, for some restricted domains—most prominently, voter interval preferences—we provide polynomial-time algorithms for finding committees that provide a good approximation to IR. Our experimental results suggest that IR is achievable in many instances that follow somewhat realistic preferences. It remains an open problem to find intuitive voting rules that provide (approximate) IR whenever possible.

# Chapter 6

# Effectiveness of Polls for Predicting Ranked Election Outcomes

## 6.1 Introduction

The aim of voting is to make a good collective decision for a group of individuals based on the preferences of its members over a set of alternatives. In the vast literature published on voting since the work of Condorcet [1785], numerous voting rules have been proposed which intuitively provide different notions of what makes an outcome best for the group. Additionally, the literature offers several frameworks that define which collective decisions are good, and allow evaluating as well as systematically designing voting rules; examples include the axiomatic approach [Arrow, 1951, Moulin, 1988], distance rationalizability [Meskanen and Nurmi, 2008, Elkind et al., 2015], noisy voting [Young, 1988, Caragiannis et al., 2016], and implicit utilitarian voting [Procaccia and Rosenschein, 2006, Boutilier et al., 2015]. However, most of this literature assumes that the voting rule in question is able to observe the preferences of every individual in the group.

In most real-world applications, only a fraction of the members actually participate in voting. The goal of the decision-making system is then to *predict* what the right collective decision is for the whole group (including individuals whose preferences are not observed) given the preferences of some of its members. In this work, we assume that there are $n$ individuals (a.k.a. voters), and we are given a voting rule $f_1$ which can make the desired collective decision given the preferences of all $n$ voters. Instead, we observe the preferences of only $k$ out of $n$ voters, and our goal is to design a voting rule $f_2$ which, when applied on the $k$ observed preferences, predicts the outcome of $f_1$ on all $n$ preferences.

This framework has two immediate motivations. As described so far, we could think of $f_1$ as the idealized voting rule we would like to implement if every individual votes, and $f_2$ as the voting rule we should really implement if only a fraction of individuals are expected to vote. Alternatively, we can imagine a setting where voting rule $f_1$ is implemented for an upcoming election, and we would like to conduct a poll to observe a subset of preferences and use $f_2$ to predict the outcome of the upcoming election. Predicting outcomes of political elections using surveys has been extensively studied [Tumasjan et al., 2010, Rothschild and Wolfers, 2011, Walther, 2015, Graefe, 2014, Lewis-Beck and Tien, 1999, Lewis-Beck, 2005, Dey and Bhattacharyya, 2015].

One question still lingers: *Which $k$ out of $n$ voters would participate?* If they are adversarially

chosen, it is not difficult to see that predicting the election outcome is impossible. For instance, if $49$ out of $100$ voters report that they prefer alternative $a$ over alternative $b$, it is impossible to know if this is in fact the majority opinion. In the worst case, each of the remaining $51$ voters might prefer $b$ over $a$. However, we can argue that this is an unlikely scenario.

In this work, we assume that the $k$ voters are sampled from a distribution.[1] This raises a number of questions: Which voting rule $f_2$, given the $k$ observed votes, would best predict the outcome of $f_1$ on all $n$ votes? When is it optimal to simply apply the same voting rule on the observed votes (i.e. $f_2 = f_1$)? Which voting rules $f_1$ can be predicted well? If we have stochastic information about the unobserved votes, how do we incorporate it into the prediction rule $f_2$? Such questions are the focus of this work.

## Our Results

We study a setting with $m$ alternatives and $n$ voters who have ranked preferences over the alternatives. A voting rule takes as input a set of rankings and returns a societal ranking. In case of ties, it may return a set of societal rankings.

We use $\vec{\sigma}_n$ to denote the profile of all $n$ ranked votes, and $\vec{\pi}_k$ to denote the sample of $k$ votes. For simplicity, assume for now that the $k$ votes are selected uniformly at random. We focus on the family of positional scoring rules, which includes popular voting rules such as plurality, Borda count, harmonic rule, $k$-approval, and veto. Given positional scoring rules $f_1$ and $f_2$, we are interested in the probability that $f_2(\vec{\pi}_k) \subseteq f_1(\vec{\sigma}_n)$, i.e., that $f_2$ predicts the outcome of $f_1$ on $\vec{\sigma}_n$ by producing a refinement of this outcome.

This probability depends on the underlying profile $\vec{\sigma}_n$. In Section 6.3, we consider the prediction accuracy in the worst case over $\vec{\sigma}_n$. We show that for a positional scoring rule $f_1$ other than plurality and veto, no positional scoring rule $f_2$ can predict its outcome with a positive probability in the worst case, when $n$ and $k$ have different parity (Theorems 26 and 27). This holds for any distribution from which the $k$ observed votes are sampled. When $f_1$ is plurality or veto, and the distribution of samples is uniform, we show that $f_1$ is the optimal predictor of itself among all positional scoring rules (Theorem 28), but its prediction accuracy is still small when the number of alternatives is large (Theorems 29 and 30). In summary, it is impossible to predict the outcome of any positional scoring rule with a reasonable accuracy when no additional information is known about $\vec{\sigma}_n$.[2]

In Section 6.4, we consider the expected prediction accuracy when $\vec{\sigma}_n$ is drawn from a known prior. Using the simplest case of two alternatives, where our goal is to predict the majority rule $f_1$, we show that the knowledge of prior can have little to significant effect on the optimal prediction rule $f_2$, depending on how large $k$ is compared to $n$ and how concentrated the prior is (Theorem 31).

Our experiments in Section 6.5 show that when $\vec{\sigma}_n$ is drawn from a concentrated prior (the Mallows model with $\varphi = 1/3$), most voting rules can be predicted with at least $98\%$ accuracy given only $3\%$ of the votes. However, when $\vec{\sigma}_n$ is drawn from the uniform prior, most voting rules cannot be predicted with accuracy more than $4\%$ given only $3\%$ of the votes, although the accuracy increases with more

[1] Our main results hold for all possible distributions, but we provide additional results for the special case where the $k$ voters are sampled uniformly at random from all $n$ voters.

[2] This is partly due to the fact that we want to predict the entire ranking of alternatives returned by the rule. In Section 6.7, we consider the weaker requirement of predicting only the winning alternative, and present a mix of positive and negative results.

observed votes. We also curiously discover that, in certain settings, the harmonic rule predicts other voting rules better than they predict themselves.

### Related Work

Most closely related to ours is the work of Dey and Bhattacharyya [2015]. They consider voting rules which output a single alternative instead of a ranking, and study the problem of predicting the output of a given voting rule on an unknown election by sampling votes. Their work differs from ours in two key aspects. First, they sample with replacement and allow the prediction rule to determine how many votes to sample. In contrast, our sampling is without replacement (which becomes dramatically different when $k$ is comparable to $n$) and the sampled votes are given. But more importantly, they assume that the underlying election has a margin of victory that is at least a constant fraction of $n$, that is, the underlying election is such that changing a constant fraction of the votes cannot change the outcome of the voting rule. We do not make this assumption. In fact, our negative results are derived precisely by considering elections that are borderline. In that sense, our results complement the results of Dey and Bhattacharyya [2015] by showing that their positive results are replaced by strong negative results when their margin of victory assumption is dropped.

Our results also have a surprising connection to the work of Borodin et al. [2019]. They consider an implicit utilitarian voting framework, in which voters and alternatives are embedded in an underlying metric space, each voter ranks the alternatives, and the social cost of an alternative is measured by its total distance from the voters. One of their results (informally) shows that given a voting rule and an *arbitrary* set of $k$ votes, where $k = \Theta(n)$, it is possible to produce an alternative that is almost as good as the alternative that would be produced by the voting rule with all $n$ votes. That is, in their framework, it is possible to do almost as well as the idealized voting rule even if the sampled votes are adversarially. This is fundamentally impossible in our setting.

A bit further afield, there is also work on predicting election outcomes under different types of uncertainty such as partial preferences [Doucette et al., 2014, Baumeister et al., 2012, Lang et al., 2012, Aziz et al., 2015], uncertainty about which voters or candidates would participate in the election (even if all preferences are known upfront) [Wojtas and Faliszewski, 2012], or distributional uncertainty about *each* voter's preferences [Hazon et al., 2012].

## 6.2 Preliminaries

For $k \in \mathbb{N}$, let $[k] = \{1, \ldots, k\}$. We consider a set $A = \{a_1, \ldots, a_m\}$ of $m$ *alternatives* and a set $N = \{1, \ldots, n\}$ of $n$ *voters*. We denote by $\mathcal{L}(A)$ the set of all rankings over $A$. We use $a \succ_\sigma b$ to denote that alternative $a$ is preferred to alternative $b$ under ranking $\sigma$. Each voter $i$ has a preference ranking (*vote*), denoted $\sigma_i \in \mathcal{L}(A)$. The *(preference) profile* $\vec{\sigma}_n = (\sigma_1, \ldots, \sigma_n)$ is the collection of all $n$ votes.

A voting rule (technically, a social welfare function) is a function $f : \mathcal{L}(A)^n \to 2^{\mathcal{L}(A)}$, which takes as input a profile and outputs a set of tied rankings. In this work, we focus on the family of *positional scoring rules*, denoted $\mathcal{F}$. A positional scoring rule $f_{\vec{s}}$ is characterized by a scoring vector $\vec{s} = (s_1, \ldots, s_m) \in \mathbb{R}^m$, where $s_t \geqslant s_{t+1}$ for each $t \in [m-1]$ and $s_1 > s_m$. Given a profile $\vec{\sigma}_n$, $f_{\vec{s}}$ assigns $s_t$ points to the $t^{\text{th}}$ alternative in voter $i$'s vote, for each $i \in N$ and $t \in [m]$. Let $\mathrm{sc}_{\vec{s}}(a, \vec{\sigma}_n) = \sum_{i=1}^{n} s_{\sigma_i(a)}$

denote the total score of $a \in A$, where $\sigma_i(a)$ is the rank of $a$ in voter $i$'s vote. Then, $f_{\vec{s}}$ returns the set of rankings where the alternatives are sorted in a non-ascending order of their scores.

We partition $\mathcal{F}$ into three subfamilies, $\mathcal{F}_1$, $\mathcal{F}_2$ and $\mathcal{F}_3$. Family $\mathcal{F}_1$ consists of all rules $f_{\vec{s}}$ for which $s_2 > s_{m-1}$. This includes the well known Borda rule ($\vec{s} = (m, m-1, \ldots, 1)$) and harmonic rule ($\vec{s} = (1, 1/2, \ldots, 1/m)$). The remaining rules $f_{\vec{s}}$ satisfy $s_2 = s_3 = \ldots = s_{m-1}$. Among these, family $\mathcal{F}_2$ consists of rules for which $s_1 > s_2 = \ldots = s_{m-1} > s_m$, while family $\mathcal{F}_3$ contains the two remaining rules: $s_1 > s_2 = \ldots = s_m$ is equivalent to plurality, and $s_1 = \ldots = s_{m-1} > s_m$ is equivalent to veto.

The *Mallows model* is a distribution over $\mathcal{L}(A)$, parametrized by a *central ranking* $\sigma^* \in \mathcal{L}(A)$ and a *noise parameter* $\varphi \in [0, 1]$. To obtain a sample ranking from this distribution, one generates an independent comparison between each pair of alternatives which matches with $\sigma^*$ with probability $p$ (where $p \geqslant 1/2$ and $\varphi = (1-p)/p$), and restarts if the comparisons violate transitivity. When $\varphi = 0$ (i.e. $p = 1$), the distribution puts all the probability mass on $\sigma^*$. When $\varphi = 1$ (i.e. $p = 1/2$), we obtain the uniform distribution, also known as impartial culture.

## 6.3 Worst-Case Predictability

Given a positional scoring rule $f_1 \in \mathcal{F}$, our goal in this chapter is to study how accurately one can predict its outcome on a profile $\vec{\sigma}_n$ given a sample of $k$ votes from the profile, where $k \leqslant n$. Specifically, let $\mathcal{S}_k(\vec{\sigma}_n)$ denote the set of all subsets of $\vec{\sigma}_n$ of size $k$, and let $U_k(\vec{\sigma}_n)$ be the uniform distribution over $\mathcal{S}_k(\vec{\sigma}_n)$. Define the accuracy of predicting $f_1$ using a positional scoring rule $f_2 \in \mathcal{F}$ on profile $\vec{\sigma}_n$ as $\mathrm{acc}(f_1, f_2, \vec{\sigma}_n) = \Pr_{\vec{\pi}_k \sim U_k(\vec{\sigma}_n)}[f_2(\vec{\pi}_k) \subseteq f_1(\vec{\sigma}_n)]$.

Note that $f_2(\vec{\pi}_k) \subseteq f_1(\vec{\sigma}_n)$ allows $f_2$ to break some of the ties produced by $f_1$ on $\vec{\sigma}_n$. This makes our negative results stronger than if we had required $f_2(\vec{\pi}_k) = f_1(\vec{\sigma}_n)$. Similarly, although we defined accuracy for $\vec{\pi}_k$ sampled from the uniform distribution $U_k(\vec{\sigma}_n)$, our main negative results (Theorems 26 and 27) hold for all distributions since they establish zero accuracy.

We define the *worst-case accuracy* of predicting $f_1$ using $f_2$ as $\mathrm{acc}(f_1, f_2) = \min_{\vec{\sigma}_n} \mathrm{acc}(f_1, f_2, \vec{\sigma}_n)$. Then, we define the *worst-case predictability* of $f_1$ as $\mathrm{acc}(f_1) = \sup_{f_2 \in \mathcal{F}} \mathrm{acc}(f_1, f_2)$, which is the worst-case accuracy of predicting $f_1$ using the best positional scoring rule $f_2$. Note that these quantities depend on $n$, $m$, and $k$, which are fixed in our framework. Motivated by political applications, we are interested in cases where $n$ is large but $k$ and $m$ are relatively smaller.

### Predicting a Rule from $\mathcal{F}_1$ or $\mathcal{F}_2$

We begin by establishing a strong negative result: every positional scoring rule in $\mathcal{F}_1$ and $\mathcal{F}_2$ has zero worst-case predictability. That is, such a rule cannot be predicted by any positional scoring rule with positive worst-case accuracy. The strength of the result lies in two observations. First, as we argued above, zero predictability implies that the outcome of the rule cannot be predicted given *any* subset of $k$ votes; thus, the negative result holds for any distribution from which the $k$ observed votes are drawn. Second, while the impossibility of prediction may be intuitive for small values of $k$, the result holds even when $k = n - 1$, i.e., when all but one of the votes are observed.

**Theorem 26.** *Let $n \geqslant 2$, $m \geqslant 7$, and $k \in [n-1]$ such that $n$ and $k$ have different parity. Then, for any positional scoring rules $f_1 \in \mathcal{F}_1$ and $f_2 \in \mathcal{F}$, we have* $\mathrm{acc}(f_1, f_2) = 0$.

*Proof.* Fix positional scoring rules $f_1 \in \mathcal{F}_1$ and $f_2 \in \mathcal{F}$. Let $\vec{r}$ and $\vec{s}$ denote their scoring vectors, respectively. Because $f_1 \in \mathcal{F}_1$, we have $r_2 > r_{m-1}$. We consider cases of even and odd $k$, and for each case, construct a profile on which $f_2$ predicts $f_1$ with zero accuracy.

**Odd $n$, even $k$:** We start with the case where $n$ is odd and $k$ is even. Consider the following profile $\vec{\sigma}_n$. Each row represents a ranking where alternatives are listed from left to right in the most preferred to least preferred order. The first two rankings appear $(n-1)/2$ times each, and the third ranking appears once. Alternatives not shown appear in an arbitrary order in the middle.

| | |
|---|---|
| $\frac{n-1}{2}$ votes | $a_1 \succ a_2 \succ a_3 \succ \ldots \succ a_{m-2} \succ a_{m-1} \succ a_m$ |
| $\frac{n-1}{2}$ votes | $a_m \succ a_{m-1} \succ a_{m-2} \succ \ldots \succ a_3 \succ a_2 \succ a_1$ |
| 1 vote | $a_1 \succ a_{m-1} \succ a_3 \succ \ldots \succ a_{m-2} \succ a_2 \succ a_m$ |

We denote with $\sigma_1$, $\sigma_2$ and $\sigma_3$ the ranking of the first, second and third rows, respectively. Because $f_1 \in \mathcal{F}_1$, it holds that for every $\sigma^* \in f_1(\vec{\sigma}_n)$, $a_1 \succ_{\sigma^*} a_m$ and $a_{m-1} \succ_{\sigma^*} a_2$. We show that for every sample $\vec{\pi}_k \in \mathcal{S}_k(\vec{\sigma}_n)$, $f_2(\vec{\pi}_k) \not\subseteq f_1(\vec{\sigma}_n)$, i.e., for some $\widehat{\sigma} \in f_2(\vec{\pi}_k)$ at least one of $a_1 \succ_{\widehat{\sigma}} a_m$ and $a_{m-1} \succ_{\widehat{\sigma}} a_2$ fails to hold. Suppose for contradiction that there exists $\vec{\pi}_k$ for which this does not happen.

Let $x_1$, $x_2$, and $x_3$ denote the number of times $\sigma_1$, $\sigma_2$, and $\sigma_3$ appear in $\vec{\pi}_k$, respectively. Note that $x_3 \in \{0, 1\}$. To have $a_1 \succ_{\widehat{\sigma}} a_m$ for every $\widehat{\sigma} \in f_2(\vec{\pi}_k)$, we need

$$\begin{aligned} &x_1 \cdot s_1 + x_2 \cdot s_m + x_3 \cdot s_1 > x_1 \cdot s_m + x_2 \cdot s_1 + x_3 \cdot s_m \\ &\Rightarrow (x_1 + x_3) \cdot (s_1 - s_m) > x_2 \cdot (s_1 - s_m). \end{aligned}$$

Given that $s_1 > s_m$, this implies $x_1 + x_3 > x_2$.

On the other hand, to have $a_{m-1} \succ_{\widehat{\sigma}} a_2$ for every $\widehat{\sigma} \in f_2(\vec{\pi}_k)$, we need

$$\begin{aligned} &x_1 \cdot s_{m-1} + x_2 \cdot s_2 + x_3 \cdot s_2 > x_1 \cdot s_2 + x_2 \cdot s_{m-1} + x_3 \cdot s_{m-1} \\ &\Rightarrow (x_2 + x_3) \cdot (s_2 - s_{m-1}) > x_1 \cdot (s_2 - s_{m-1}). \end{aligned}$$

Given $s_2 \geqslant s_{m-1}$, this implies $x_2 + x_3 > x_1$.

Given $x_1 + x_3 > x_2$ and $x_2 + x_3 > x_1$, we can derive $x_3 > 0$, i.e., $x_3 = 1$. But then, we have $x_1 + 1 > x_2$ and $x_2 + 1 > x_1$, which implies $x_1 = x_2$. In this case, $|\vec{\pi}_k| = x_1 + x_2 + 1$ is odd, which contradicts the fact that $k$ is even.

**Even $n$, odd $k$:** We now consider the case of even $n$ and odd $k$. We begin by establishing the following property of $f_1$. Recall that for the scoring vector $\vec{r}$ of $f_1 \in \mathcal{F}_1$, we have $r_2 > r_{m-1}$.

**Lemma 19.** *There exists a $t \in \{3, \ldots, m-3\}$ such that $r_2 - r_{m-1} > r_t - r_{t+1}$.*

*Proof.* First, suppose there exists a $p \in \{2, \ldots, m-2\}$ such that $r_2 = r_p > r_{p+1} = r_{m-1}$ (i.e., in going from $r_2$ to $r_{m-1}$, the score drops only once). If $p \geqslant 4$, then we set $t = 3$. In this case, we have $r_t - r_{t+1} = 0 < r_2 - r_{m-1}$, as desired. If $p \leqslant 3$, then we set $t = 4$. Because $m \geqslant 7$, we have $t \leqslant m - 3$. Also, we again have $r_t - r_{t+1} = 0 < r_2 - r_{m-1}$.

Next, suppose there exist distinct $p, q \in \{2, \ldots, m-2\}$ such that $r_p > r_{p+1}$ and $r_q > r_{q+1}$ (i.e., in going from $r_2$ to $r_{m-1}$, the score drops at least twice). Then, we can simply set $t = 3$. This ensures that $r_t - r_{t+1} < r_2 - r_{m-1}$ (if it were equal, then the score would drop only once in going from $r_2$ to $r_{m-1}$). This completes the proof. □

Let us fix $t \in \{3, \ldots, m-3\}$ for which Lemma 19 holds. Consider the following profile $\vec{\sigma}_n$.

| | |
|---|---|
| $\frac{n-2}{2}$ votes | $a_1 \succ a_2 \succ a_3 \succ a_4 \succ \ldots \succ a_{m-2} \succ a_{m-1} \succ a_m$ |
| $\frac{n-2}{2}$ votes | $a_m \succ a_{m-1} \succ a_{m-2} \succ \ldots \succ a_4 \succ a_3 \succ a_2 \succ a_1$ |
| 1 vote | $a_1 \succ a_3 \succ \ldots \succ a_2 \succ a_{m-1} \succ \ldots \succ a_{m-2} \succ a_m$ |
| 1 vote | $a_3 \succ a_{m-1} \succ \ldots \succ a_m \succ a_1 \succ \ldots \succ a_2 \succ a_{m-2}$ |

We again denote with $\sigma_1$, $\sigma_2$, $\sigma_3$ and $\sigma_4$ the rankings in rows 1, 2, 3 and 4, respectively. In ranking $\sigma_3$, $a_2$ is at position $t$ and $a_{m-1}$ is at position $t+1$. In the ranking $\sigma_4$, $a_m$ is at position $t$ and $a_1$ is at position $t+1$. In each ranking, alternatives not shown appear in the unfilled positions arbitrarily.

First, we argue about the outcome of $f_1$ on this profile. From the Lemma 19, it is obvious that $r_1 - r_m \geqslant r_2 - r_{m-1} > r_t - r_{t+1}$. Using this, it is easy to see that for every $\sigma^* \in f_1(\vec{\sigma}_n)$, $a_1 \succ_{\sigma^*} a_m$ and $a_{m-1} \succ_{\sigma^*} a_2$. We now argue that for every sample $\vec{\pi}_k \in \mathcal{S}_k(\vec{\sigma}_n)$, there exists $\widehat{\sigma} \in f_2(\vec{\pi}_k)$ which violates at least one of $a_1 \succ_{\widehat{\sigma}} a_m$ and $a_{m-1} \succ_{\widehat{\sigma}} a_2$. Suppose for contradiction that there exists a sample $\vec{\pi}_k$ for which this does not happen.

Again, let $x_1$, $x_2$, $x_3$, and $x_4$ denote the number of times $\sigma_1$, $\sigma_2$, $\sigma_3$, and $\sigma_4$ appear in $\vec{\pi}_k$, respectively. Note that $x_3, x_4 \in \{0, 1\}$. To have $a_1 \succ_{\widehat{\sigma}} a_m$ for every $\widehat{\sigma} \in f_2(\vec{\pi}_k)$, we need

$$\begin{aligned} &x_1 \cdot s_1 + x_2 \cdot s_m + x_3 \cdot s_1 + x_4 \cdot s_{t+1} > x_1 \cdot s_m + x_2 \cdot s_1 + x_3 \cdot s_m + x_4 \cdot s_t \\ &\Rightarrow (x_1 + x_3) \cdot (s_1 - s_m) > x_2 \cdot (s_1 - s_m) + x_4 \cdot (s_t - s_{t+1}). \end{aligned} \tag{6.1}$$

On the other hand, to have $a_{m-1} \succ_{\widehat{\sigma}} a_2$ for every $\widehat{\sigma} \in f_2(\vec{\pi}_k)$, we need

$$\begin{aligned} &x_1 \cdot s_{m-1} + x_2 \cdot s_2 + x_3 \cdot s_{t+1} + x_4 \cdot s_2 > x_1 \cdot s_2 + x_2 \cdot s_{m-1} + x_3 \cdot s_t + x_4 \cdot s_{m-1} \\ &\Rightarrow (x_2 + x_4) \cdot (s_2 - s_{m-1}) > x_1 \cdot (s_2 - s_{m-1}) + x_3 \cdot (s_t - s_{t+1}). \end{aligned} \tag{6.2}$$

We now distinguish between four cases.

**Case 1:** $x_3 = x_4 = 0$**.** Since $s_1 > s_m$, from Equation (6.1), we obtain $x_1 > x_2$. Moreover, as $s_2 \geqslant s_{m-1}$, from Equation (6.2), we obtain $x_2 > x_1$, which is a contradiction.

**Case 2:** $x_3 = 1$ **and** $x_4 = 0$**.** Since $s_1 > s_m$, from Equation (6.1), we obtain $x_1 + 1 > x_2$. Moreover, as $s_2 \geqslant s_{m-1}$ and $s_t \geqslant s_{t+1}$, from Equation (6.2), we obtain $x_2 > x_1$, which is a contradiction.

**Case 3:** $x_3 = 0$ **and** $x_4 = 1$**.** This case leads to a contradiction in a manner similar to Case 2, so we omit the details.

**Case 4:** $x_3 = 1$ **and** $x_4 = 1$**.** Since $s_1 > s_m$ and $s_t \geqslant s_{t+1}$, from Equation (6.1) we obtain that $x_1 + 1 > x_2$. Similarly, since $s_2 \geqslant s_{m-1}$ and $s_t \geqslant s_{t+1}$, from Equation (6.2) we obtain $x_2 + 1 > x_1$. This implies $x_1 = x_2$, which implies $|\vec{\pi}_k| = x_1 + x_2 + 2$ is even, which in turn contradicts the fact that $k$ is odd. □

Unfortunately, the proof of Theorem 26 does not directly work when $f_1 \in \mathcal{F}_2$. However, a similar proof with somewhat more intricate profiles works, yielding the following result.

**Theorem 27.** *Let $n \geqslant 4$, $m \geqslant 5$, and $k \in [n-1]$ such that $n$ and $k$ have different parity. Then, for any positional scoring rules $f_1 \in \mathcal{F}_2$ and $f_2 \in \mathcal{F}$, we have* $\mathrm{acc}(f_1, f_2) = 0$.

*Proof.* Fix positional scoring rules $f_1 \in \mathcal{F}_2$ and $f_2 \in \mathcal{F}$. Let $\vec{r}$ and $\vec{s}$ denote their scoring vectors, respectively. We consider cases of even and odd $k$, and for each case, construct a profile on which $f_2$ predicts $f_1$ with zero accuracy.

**Odd $n$, even $k$:** We start with the case where $n$ is odd and $k$ is even. Consider the following profile $\vec{\sigma}_n$. Alternatives not shown appear in an arbitrary order in the middle.

| | |
|---|---|
| $\frac{n-1}{2}$ votes | $a_1 \succ a_m \succ a_3 \succ \ldots \succ a_{m-2} \succ a_2 \succ a_{m-1}$ |
| $\frac{n-1}{2}$ votes | $a_m \succ a_1 \succ a_3 \succ \ldots \succ a_{m-2} \succ a_{m-1} \succ a_2$ |
| 1 vote | $a_1 \succ a_m \succ a_3 \succ \ldots \succ a_{m-2} \succ a_{m-1} \succ a_2$ |

We denote with $\sigma_1$, $\sigma_2$ and $\sigma_3$ the ranking of the first, second and third rows, respectively. Because $f_1 \in \mathcal{F}_2$, we have $r_1 > r_2 = r_{m-1} > r_m$, and hence for every $\sigma^* \in f_1(\vec{\sigma}_n)$, it holds $a_1 \succ_{\sigma^*} a_m$ and $a_{m-1} \succ_{\sigma^*} a_2$. We show that for every sample $\vec{\pi}_k \in \mathcal{S}_k(\vec{\sigma}_n)$, $f_2(\vec{\pi}_k) \not\subseteq f_1(\vec{\sigma}_n)$, i.e., for some $\widehat{\sigma} \in f_2(\vec{\pi}_k)$ at least one of $a_1 \succ_{\widehat{\sigma}} a_m$ and $a_{m-1} \succ_{\widehat{\sigma}} a_2$ fails to hold. Suppose for contradiction that there exists $\vec{\pi}_k$ for which this does not happen.

Let $x_1$, $x_2$, and $x_3$ denote the number of times $\sigma_1$, $\sigma_2$, and $\sigma_3$ appear in $\vec{\pi}_k$, respectively. Note that $x_3 \in \{0, 1\}$. To have $a_1 \succ_{\widehat{\sigma}} a_m$ for every $\widehat{\sigma} \in f_2(\vec{\pi}_k)$, we need

$$\begin{aligned} &x_1 \cdot s_1 + x_2 \cdot s_2 + x_3 \cdot s_1 > x_1 \cdot s_2 + x_2 \cdot s_1 + x_3 \cdot s_2. \\ &\Rightarrow (x_1 + x_3) \cdot (s_1 - s_2) > x_2 \cdot (s_1 - s_2). \end{aligned}$$

Given that $s_1 \geqslant s_2$, this implies $x_1 + x_3 > x_2$.

On the other hand, to have $a_{m-1} \succ_{\widehat{\sigma}} a_2$ for every $\widehat{\sigma} \in f_2(\vec{\pi}_k)$, we need

$$\begin{aligned} &x_1 \cdot s_m + x_2 \cdot s_{m-1} + x_3 \cdot s_{m-1} > x_1 \cdot s_{m-1} + x_2 \cdot s_m + x_3 \cdot s_m \\ &\Rightarrow (x_2 + x_3) \cdot (s_{m-1} - s_m) > x_1 \cdot (s_{m-1} - s_m). \end{aligned}$$

Given $s_{m-1} \geqslant s_m$, this implies $x_2 + x_3 > x_1$.

Given $x_1 + x_3 > x_2$ and $x_2 + x_3 > x_1$, we can derive $x_3 > 0$, i.e., $x_3 = 1$. But then, we have $x_1 + 1 > x_2$ and $x_2 + 1 > x_1$, which implies $x_1 = x_2$. In this case, $|\vec{\pi}_k| = x_1 + x_2 + 1$ is odd, which contradicts the fact that $k$ is even.

**Even $n$, odd $k$:** We now consider the case of even $n$ and odd $k$. Consider the following profile $\vec{\sigma}_n$.

| | |
|---|---|
| $\frac{n-2}{2}$ votes | $a_1 \succ a_m \succ a_3 \succ \ldots \succ a_{m-2} \succ a_2 \succ a_{m-1}$ |
| $\frac{n-2}{2}$ votes | $a_m \succ a_1 \succ a_3 \succ \ldots \succ a_{m-2} \succ a_{m-1} \succ a_2$ |
| 1 vote | $a_1 \succ a_m \succ \ldots \succ a_2 \succ a_{m-1} \succ \ldots \succ a_{m-2} \succ a_3$ |
| 1 vote | $a_3 \succ a_4 \succ \ldots \succ a_m \succ a_1 \succ \ldots \succ a_{m-1} \succ a_2$ |

We again denote with $\sigma_1$, $\sigma_2$, $\sigma_3$ and $\sigma_4$, the rankings in rows 1, 2, 3 and 4, respectively. Fix an arbitrary $t$ such that $2 < t < m - 2$. In ranking $\sigma_3$, alternative $a_2$ is at some position $t$ and alternative $a_{m-1}$ is at position $t + 1$. In ranking $\sigma_4$, alternative $a_m$ is at position $t$ and alternative $a_1$ is at position $t + 1$. In each ranking, alternatives not shown appear in the unfilled positions arbitrarily.

First, we argue about the outcome of $f_1$ on this profile. Recall that for the scoring vector $\vec{r}$ of $f_1 \in \mathcal{F}_2$, we have $r_1 > r_2$ and $r_{m-1} > r_m$, while $r_t = r_{t+1}, \forall t \in [2, m-2]$. Using this, it is easy to see that for every $\sigma^* \in f_1(\vec{\sigma}_n)$, $a_1 \succ_{\sigma^*} a_m$ and $a_{m-1} \succ_{\sigma^*} a_2$. We now argue that for every sample $\vec{\pi}_k \in \mathcal{S}_k(\vec{\sigma}_n)$, there exists $\widehat{\sigma} \in f_2(\vec{\pi}_k)$ which violates at least one of $a_1 \succ_{\widehat{\sigma}} a_m$ and $a_{m-1} \succ_{\widehat{\sigma}} a_2$. Suppose for contradiction that there exists a sample $\vec{\pi}_k$ for which this does not happen. Here, we have to take four cases depending on whether each of $\sigma_3$ and $\sigma_4$ appears in $\vec{\pi}_k$.

Again, let $x_1$, $x_2$, $x_3$, and $x_4$ denote the number of times $\sigma_1$, $\sigma_2$, $\sigma_3$, and $\sigma_4$ appear in $\vec{\pi}_k$, respectively. Note that $x_3, x_4 \in \{0, 1\}$. To have $a_1 \succ_{\widehat{\sigma}} a_m$ for every $\widehat{\sigma} \in f_2(\vec{\pi}_k)$, we need

$$\begin{aligned} &x_1 \cdot s_1 + x_2 \cdot s_2 + x_3 \cdot s_1 + x_4 \cdot s_{t+1} > x_1 \cdot s_2 + x_2 \cdot s_1 + x_3 \cdot s_2 + x_4 \cdot s_t \\ &\Rightarrow (x_1 + x_3) \cdot (s_1 - s_2) > x_2 \cdot (s_1 - s_2) + x_4 \cdot (s_t - s_{t+1}). \end{aligned} \tag{6.3}$$

On the other hand, to have $a_{m-1} \succ_{\widehat{\sigma}} a_2$ for every $\widehat{\sigma} \in f_2(\vec{\pi}_k)$, we need

$$\begin{aligned} &x_1 \cdot s_m + x_2 \cdot s_{m-1} + x_3 \cdot s_{t+1} + x_4 \cdot s_{m-1} \\ &\qquad > x_1 \cdot s_{m-1} + x_2 \cdot s_m + x_3 \cdot s_t + x_4 \cdot s_m \\ &\Rightarrow (x_2 + x_4) \cdot (s_{m-1} - s_m) > x_1 \cdot (s_{m-1} - s_m) + x_3 \cdot (s_t - s_{t+1}). \end{aligned} \tag{6.4}$$

We now distinguish between four cases.

**Case 1:** $x_3 = x_4 = 0$**.** Since $s_1 > s_2$, from Equation (6.3), we obtain $x_1 > x_2$. Moreover, as $s_{m-1} \geqslant s_m$, from Equation (6.4), we obtain $x_2 > x_1$, which is a contradiction.

**Case 2:** $x_3 = 1$ **and** $x_4 = 0$**.** Since $s_1 > s_2$, from Equation (6.3), we obtain $x_1 + 1 > x_2$. Moreover, as $s_{m-1} \geqslant s_m$ and $s_t \geqslant s_{t+1}$, from Equation (6.4), we obtain $x_2 > x_1$, which is a contradiction.

**Case 3:** $x_3 = 0$ **and** $x_4 = 1$**.** This case leads to a contradiction in a manner similar to Case 2, so we omit the details.

**Case 4:** $x_3 = 1$ **and** $x_4 = 1$**.** Since $s_1 > s_2$ and $s_t \geqslant s_{t+1}$, from Equation (6.3) we obtain that $x_1 + 1 > x_2$. Similarly, since $s_{m-1} \geqslant s_m$ and $s_t \geqslant s_{t+1}$, from Equation (6.4) we obtain $x_2 + 1 > x_1$. This implies $x_1 = x_2$, which implies $|\vec{\pi}_k| = x_1 + x_2 + 2$ is even, which in turn contradicts the fact that $k$ is odd.

□

## Predicting a Rule from $\mathcal{F}_3$

The remaining family $\mathcal{F}_3$ contains exactly two voting rules: plurality (denoted $f_{\text{plu}}$) and veto (denoted $f_{\text{veto}}$). These two rules are special within the family of positional scoring rules. While Theorems 26 and 27 establish that every other positional scoring rule has zero worst-case predictability, we will show that this is not the case with plurality or veto.

**Proposition 6.** *Let $k \geqslant m(m-1)/2$. Then, we have* $\mathrm{acc}(f_{plu}, f_{plu}) > 0$ *and* $\mathrm{acc}(f_{veto}, f_{veto}) > 0$.

*Proof.* We provide a proof for plurality. The proof for veto is similar. Consider any profile $\vec{\sigma}_n$. Let $x_i$ denote the number of times alternative $a_i$ is ranked first. Without loss of generality, assume $x_i \geqslant x_{i+1}$ for all $i \in [m-1]$. Now, we construct a sample $\vec{\pi}_k \in \mathcal{S}_k(\vec{\sigma}_n)$ such that $f_{\text{plu}}(\vec{\pi}_k) \subseteq f_{\text{plu}}(\vec{\sigma}_n)$.

Let $y_i = |\{j \in [i, m-1] : x_j > x_{j+1}\}|$ for each $i \in [m-1]$, and $y_m = 0$. We begin by choosing $y_i$ arbitrary rankings from $\vec{\sigma}_n$ which rank $a_i$ first, for each $i$, and adding them to the sample. It is easy to see that $y_i = y_{i+1}$ if and only if $x_i = x_{i+1}$ and $y_i > y_{i+1}$ if and only if $x_i > x_{i+1}$, for all $i \in [m-1]$. Further, $\sum_{i=1}^{m} y_i \leqslant m(m-1)/2 \leqslant k$. If we ran plurality on the sample constructed so far, the set of rankings it returns would be precisely $f_{\text{plu}}(\vec{\sigma}_n)$. However, this sample may contain fewer than $k$ votes.

We complete the sample $\vec{\pi}_k$ by adding any remaining votes which rank $a_1$ first, then adding any remaining votes which rank $a_2$ first, etc, until the sample size becomes $k$. Let $z_i$ denote the final number of votes in the sample which rank $a_i$ first. Then, for all $i, j \in [m]$, $x_i > x_j$ implies $y_i > y_j$, which implies $z_i > z_j$. Thus, $f_{\text{plu}}(\vec{\pi}_k) \subseteq f_{\text{plu}}(\vec{\sigma}_n)$. □

Proposition 6 raises two important questions: a) *How well can plurality or veto predict itself?*; and b) *Can some positional scoring rule predict plurality (resp. veto) better than plurality (resp. veto) itself?*

We begin by answering the latter question negatively. We show that among all positional scoring rules, the best predictor of plurality (resp. veto) is plurality (resp. veto) itself.

**Theorem 28.** *For every positional scoring rule $f_2 \in \mathcal{F}$, we have that* $\text{acc}(f_{plu}, f_2) \leqslant \text{acc}(f_{plu}, f_{plu})$ *and* $\text{acc}(f_{veto}, f_2) \leqslant \text{acc}(f_{veto}, f_{veto})$.

*Proof.* Fix $f_2 \in \mathcal{F}$. We start by showing that $\text{acc}(f_{\text{plu}}, f_2) \leqslant \text{acc}(f_{\text{plu}}, f_{\text{plu}})$. With similar arguments, we can show that for $\text{acc}(f_{\text{veto}}, f_2) \leqslant \text{acc}(f_{\text{veto}}, f_{\text{veto}})$.

Specifically, we show that for every profile $\vec{\sigma}_n$, there exists a profile $\vec{\tau}_n$ such that $\text{acc}(f_{\text{plu}}, f_{\text{plu}}, \vec{\sigma}_n) \geqslant \text{acc}(f_{\text{plu}}, f_2, \vec{\tau}_n)$. This implies the desired result.

Consider any profile $\vec{\sigma}_n$. Fix $\sigma^* \in f_{\text{plu}}(\vec{\sigma}_n)$. We construct the profile $\vec{\tau}_n$ as follows. In each ranking $\tau_i$, the alternative ranked first in $\sigma_i$ is also ranked first, and the remaining alternatives are in the opposite order of how they appear in $\sigma^*$. Because we do not change the alternatives in the first position, we have that $\text{acc}(f_{\text{plu}}, f_{\text{plu}}, \vec{\sigma}_n) = \text{acc}(f_{\text{plu}}, f_{\text{plu}}, \vec{\tau}_n)$. We now show that $\text{acc}(f_{\text{plu}}, f_{\text{plu}}, \vec{\tau}_n) \geqslant \text{acc}(f_{\text{plu}}, f_2, \vec{\tau}_n)$. More specifically, we show that for every sample $\vec{\pi}_k \in \mathcal{S}_k(\vec{\tau}_n)$, $f_{\text{plu}}(\vec{\pi}_k) \not\subseteq f_{\text{plu}}(\vec{\tau}_n)$ implies $f_2(\vec{\pi}_k) \not\subseteq f_{\text{plu}}(\vec{\tau}_n)$.

Consider any sample $\vec{\pi}_k \in \mathcal{S}_k(\vec{\tau}_n)$. Let $x_i$ and $y_i$ denote the number of times $a_i$ is ranked first in $\vec{\tau}_n$ and $\vec{\pi}_k$, respectively. Suppose $f_{\text{plu}}(\vec{\pi}_k) \not\subseteq f_{\text{plu}}(\vec{\tau}_n)$. Then, there exist alternatives $a_i, a_j \in A$ such that $x_i > x_j$ but $y_i \leqslant y_j$. Note that $x_i > x_j$ implies that $a_i \succ_{\sigma^*} a_j$. Hence, in every ranking in $\vec{\tau}_n$ (and therefore in $\vec{\pi}_k$) where $a_i$ or $a_j$ is not ranked first, $a_j$ must appear before $a_i$ (since we order them in the opposite order of $\sigma^*$). This, together with $y_j \geqslant y_i$, implies that under $\vec{\pi}_k$, $f_2$ assigns at least as much score to $a_j$ as to $a_i$. Hence, there exists $\widehat{\sigma} \in f_2(\vec{\pi}_k)$ for which $a_j \succ_{\widehat{\sigma}} a_i$, and thus $\widehat{\sigma} \notin f_{\text{plu}}(\vec{\tau}_n)$. Hence, we conclude $f_2(\vec{\pi}_k) \not\subseteq f_{\text{plu}}(\vec{\tau}_n)$. □

From Theorem 28, an upper bound on the worst-case accuracy of predicting plurality using plurality gives us an upper bound on the worst-case predictability of plurality. The same holds for veto. While Proposition 6 shows that this quantity is non-zero for $k \geqslant m(m-1)/2$, we show that it is still exponentially small in $m$ when $k$ is small compared to $n$.

**Theorem 29.** *For $n \geqslant (m-1)(m-2)/2$ and $k \leqslant cn$, where $c < 1$ is a constant, we have* $\text{acc}(f_{plu}) = \text{acc}(f_{plu}, f_{plu}) \leqslant c^{\Omega(m^2)}$ *and* $\text{acc}(f_{veto}) = \text{acc}(f_{veto}, f_{veto}) \leqslant c^{\Omega(m^2)}$.

*Proof.* We provide a proof for plurality. The proof for veto is similar. Note that $\text{acc}(f_{\text{plu}}) = \text{acc}(f_{\text{plu}}, f_{\text{plu}})$ follows from Theorem 28. We now show that there exists a profile $\vec{\sigma}_n$ for which $\text{acc}(f_{\text{plu}}, f_{\text{plu}}) \leqslant \text{acc}(f_{\text{plu}}, f_{\text{plu}}, \vec{\sigma}_n) \leqslant c^m$.

Consider the profile $\vec{\sigma}_n$ in which alternative $a_i$ appears first in exactly $i-1$ rankings, for each $i \in [m-1]$. In every other ranking, alternative $a_m$ appears first. This is feasible because $n \geqslant (m-1)(m-2)/2$. Note that for $f_{\text{plu}}(\vec{\pi}_k) \subseteq f_{\text{plu}}(\vec{\sigma}_n)$, sample $\vec{\pi}_k$ must at least contain all of $t = (m-1)(m-2)/2$ rankings in which an alternative from $\{a_1, \ldots, a_{m-1}\}$ is ranked first. For $k < (m-1)(m-2)/2$, this happens with zero probability. For $k \geqslant (m-1)(m-2)/2$, this happens with probability at most $\binom{n-t}{k-t}/\binom{n}{k} \leqslant (k/n)^t \leqslant c^t$. □

We conjecture that even when $k = n - o(n)$, $\mathrm{acc}(f_{\text{plu}}, f_{\text{plu}})$ and $\mathrm{acc}(f_{\text{veto}}, f_{\text{veto}})$ are still $O(1/m)$. For $k = n-1$, it is easy to see that they are in fact $\Theta(1/m)$.

**Theorem 30.** *For $k = n-1$, $\mathrm{acc}(f_{plu}) = \mathrm{acc}(f_{plu}, f_{plu}) = \Theta(1/m)$ and $\mathrm{acc}(f_{veto}) = \mathrm{acc}(f_{veto}, f_{veto}) = \Theta(1/m)$.*

*Proof.* Once again, we provide a proof for plurality. The proof for veto is similar. Assume $n \gg m$. Given Theorem 28, we simply need to show that $\mathrm{acc}(f_{\text{plu}}, f_{\text{plu}}) = \Theta(1/m)$.

First, we show the upper bound. For all $i \in [m]$, define

$$x_i = \begin{cases} \lceil n/m \rceil + 1 & \text{if } i = 1 \\ \lceil n/m \rceil & \text{if } 2 \leqslant i \leqslant n \mod m \\ \lfloor n/m \rfloor & \text{if } n \mod m \leqslant i \leqslant m-1 \\ \lfloor n/m \rfloor - 1 & \text{if } i = m \end{cases}$$

Note that this satisfies $x_m < x_i$ for all $i \in [m-1]$ and $\sum_{i=1}^{m} x_i = n$. Now, consider a profile $\vec{\sigma}_n$ in which alternative $a_i$ is in the top position in $x_i$ rankings, for each $i \in [m]$. For every $\sigma^* \in f_{\text{plu}}(\vec{\sigma}_n)$, we have $a_i \succ_{\sigma^*} a_m$ for all $i \in [m-1]$. Consider a sample $\vec{\pi}_{n-1} \in \mathcal{S}_{n-1}(\vec{\sigma}_n)$. To have $a_i \succ_{\widehat{\sigma}} a_m$ for all $i \in [m-1]$ and $\widehat{\sigma} \in f_{\text{plu}}(\vec{\pi}_{n-1})$, $\vec{\pi}_{n-1}$ must contain all rankings of $\vec{\sigma}_n$ except a ranking in which $a_m$ appears first. This happens with probability $(\lfloor n/m \rfloor - 1)/n = O(1/m)$. Hence, $\mathrm{acc}(f_{\text{plu}}, f_{\text{plu}}) \leqslant \mathrm{acc}(f_{\text{plu}}, f_{\text{plu}}, \vec{\sigma}_n) = O(1/m)$.

Next, we show the lower bound. Consider any profile $\vec{\sigma}_n$. For $i \in [m]$, let $x_i$ denote the number of times alternative $a_i$ appears first. Without loss of generality, assume $x_i \geqslant x_{i+1}$ for $i \in [m-1]$. Let $i^*$ be the smallest index such that $x_{i^*} \geqslant x_{i^*+1} + 2$ (if $x_i \leqslant x_{i+1} + 1$ for all $i \in [m-1]$, then let $i^* = m$). It is easy to see that $x_{i^*} = \Omega(n/m)$, and for any $\vec{\pi}_{n-1} \in \mathcal{S}_{n-1}(\vec{\sigma}_n)$ which is obtained by removing one of the rankings in which $a_{i^*}$ appears first, $f_{\text{plu}}(\vec{\pi}_{n-1}) \subseteq f_{\text{plu}}(\vec{\sigma}_n)$. Hence, $\mathrm{acc}(f_{\text{plu}}, f_{\text{plu}}) = \Omega(n/m)/n = \Omega(1/m)$. □

## 6.4 Average-Case Predictability

In the previous section, we considered the accuracy of predicting the outcome of a voting rule $f_1$ using a voting rule $f_2$ in the *worst case* over the underlying profile $\vec{\sigma}_n$, and defined $\mathrm{acc}(f_1, f_2) = \min_{\vec{\sigma}_n} \mathrm{acc}(f_1, f_2, \sigma_n)$.

In this section, we take a less pessimistic viewpoint, assume that the profile $\vec{\sigma}_n$ consists of $n$ rankings drawn iid from a known prior $\mathcal{D}$, and define $\mathrm{acc}^{\mathcal{D}}(f_1, f_2) = \mathsf{E}_{\vec{\sigma}_n \sim \mathcal{D}^n}[\mathrm{acc}(f_1, f_2, \vec{\sigma}_n)]$, where $\vec{\sigma}_n \sim \mathcal{D}^n$ denotes that $\vec{\sigma}_n$ is drawn from the product distribution $\mathcal{D}^n$.

We show that this leads to interesting phenomena even in the simplest setting with two alternatives. Let $A = \{a, b\}$. Without loss of generality, suppose $\mathcal{D}$ generates $a \succ b$ with probability $p \geqslant 1/2$ and $b \succ a$ with probability $1 - p$. This coincides with the Mallows model with central ranking $\sigma^* = a \succ b$ and noise parameter $\varphi = (1 - p)/p \in [0, 1]$.

For two alternatives, all reasonable voting rules (including all positional scoring rules) coincide with plurality, which is simply the majority rule. Hence, we fix the target voting rule as plurality ($f_1 = f_{\text{plu}}$). Our goal is to predict which of $a \succ b$ and $b \succ a$ appears more frequently in the underlying profile $\vec{\sigma}_n$. Without any distributional information, we cannot outperform running plurality on the sample $\vec{\pi}_k$, i.e., using $f_2 = f_{\text{plu}}$ (Theorem 28). However, with the knowledge of the prior, the optimal rule $f_2$ which maximizes $\text{acc}^{\mathcal{D}}(f_1, f_2)$ computes the posterior distribution of $\vec{\sigma}_n$ given both sample $\vec{\pi}_k$ and prior $\mathcal{D}$, and returns the more likely outcome of plurality on $\vec{\sigma}_n$ drawn from the posterior.

When the sample contains at least as many $a \succ b$ as $b \succ a$, the optimal rule would also return $a \succ b$. However, when the sample contains more $b \succ a$ than $a \succ b$, there is tension between the sample and the prior, and the output of the optimal rule is less clear.

Consider the extreme case in which the sample $\vec{\pi}_k$ consists of $k$ copies of $b \succ a$. If $k \geqslant n/2$, the optimal rule safely returns $b \succ a$. When $k < n/2$, the optimal rule returns $b \succ a$ if $\Pr[f_{\text{plu}}(\vec{\sigma}_n) = b \succ a | \vec{\pi}_k] > \Pr[f_{\text{plu}}(\vec{\sigma}_n) = a \succ b | \vec{\pi}_k]$, but returns $a \succ b$ otherwise. It is easy to show that $\Pr[f_{\text{plu}}(\vec{\sigma}_n) = a \succ b | \vec{\pi}_k]$ is monotonically decreasing in $k$ and in $\varphi$. Hence, there exists a unique $\varphi_k^*$ such that the optimal rule returns $a \succ b$ when $\varphi < \varphi_k^*$ and returns $b \succ a$ when $\varphi > \varphi_k^*$. Further, $\varphi_k^*$ is monotonically decreasing in $k$. The next result sheds more light on the relation between $\varphi_k^*$ and $k$. Its proof is provided in the full version.

**Theorem 31.** *Let $n \geqslant 5$ and $n - 1$ be divisible by 4. Given a sample $\vec{\pi}_k$ which consists of $k$ copies of $b \succ a$, let $\varphi_k^*$ be such that the optimal predictor returns $a \succ b$ if $\varphi < \varphi_k^*$ and $b \succ a$ if $\varphi > \varphi_k^*$. Then the following hold.*

1. *For $k = 1$, $\varphi_k^* \geqslant 1 - \frac{4 \ln n}{n+1}$.*
2. *For $k = (n-1)/2$, $\varphi_k^* \leqslant \frac{2}{n+1}$.*
3. *For $k = (n-1)/4$, $\varphi_k^* \in [1/4, 2/3]$.*

*Proof.* We begin by examining the probability of each possible underlying plurality outcome given the samples. Recall that under the prior, the probability of generating $a \succ b$ is $p \in [1/2, 1]$, and we have $\varphi = (1 - p)/p$. Given a sample $\vec{\pi}_k$ which consists of $k$ copies of $b \succ a$, we get

$$
\begin{aligned}
&\Pr[f_{\text{plu}}(\vec{\sigma}_n) = a \succ b | \vec{\pi}_k] \\
&= \frac{\Pr[f_{\text{plu}}(\vec{\sigma}_n) = a \succ b \wedge \vec{\pi}_k]}{\Pr[\vec{\pi}_k]} \\
&= \frac{\sum_{\vec{\sigma}_n} \Pr[\vec{\sigma}_n] \cdot \Pr[f_{\text{plu}}(\vec{\sigma}_n) = a \succ b \wedge \vec{\pi}_k | \vec{\sigma}_n]}{\Pr[\vec{\pi}_k]} \\
&= \frac{\sum_{\vec{\sigma}_n} \Pr[\vec{\sigma}_n] \cdot \Pr[f_{\text{plu}}(\vec{\sigma}_n)) = a \succ b | \vec{\sigma}_n] \cdot \Pr[\vec{\pi}_k | \vec{\sigma}_n]}{\Pr[\vec{\pi}_k]} \\
&= \frac{\sum_{\vec{\sigma}_n : f_{\text{plu}}(\vec{\sigma}_n) = a \succ b} \Pr[\vec{\sigma}_n] \cdot \Pr[\vec{\pi}_k | \vec{\sigma}_n]}{\Pr[\vec{\pi}_k]}
\end{aligned}
$$

$$= \frac{\sum_{i=k}^{\frac{n-1}{2}} \binom{n}{i} p^{n-i}(1-p)^i \frac{\binom{i}{k}}{\binom{n}{k}}}{\Pr[\vec{\pi}_k]}$$
$$= \frac{\sum_{i=k}^{\frac{n-1}{2}} \binom{i}{k}\binom{n}{i}\varphi^i p^n}{\binom{n}{k}\Pr[\vec{\pi}_k]}. \tag{6.5}$$

Similarly, we conclude that:

$$\Pr[f_{\text{plu}}(\vec{\sigma}_n) = b \succ a|\vec{\pi}_k] = = \frac{\sum_{i=\frac{n+1}{2}}^{n} \binom{i}{k}\binom{n}{i}\varphi^i p^n}{\binom{n}{k}\Pr[\vec{\pi}_k]}. \tag{6.6}$$

Recall that $\varphi_k^*$ is the unique value which satisfies $\Pr[f_{\text{plu}}(\vec{\sigma}_n) = a \succ b|\vec{\pi}_k] = \Pr[f_{\text{plu}}(\vec{\sigma}_n) = b \succ a|\vec{\pi}_k]$, i.e.,

$$\sum_{i=k}^{\frac{n-1}{2}} \binom{i}{k}\binom{n}{i}(\varphi_k^*)^i = \sum_{i=\frac{n+1}{2}}^{n} \binom{i}{k}\binom{n}{i}(\varphi_k^*)^i.$$

**Case 1:** $k = 1$**.** We want to show that $\varphi_k^* \geqslant 1 - \frac{4\ln n}{n+1}$. Hence, it is sufficient to show that for all $\varphi < 1 - \frac{4\ln n}{n+1}$,

$$\sum_{i=k}^{\frac{n-1}{2}} \binom{i}{k}\binom{n}{i}\varphi^i > \sum_{i=\frac{n+1}{2}}^{n} \binom{i}{k}\binom{n}{i}\varphi^i.$$

For $k = 1$, this reduces to

$$\sum_{i=1}^{\frac{n-1}{2}} \binom{n}{i} \cdot i \cdot \varphi^i > \sum_{i=\frac{n+1}{2}}^{n} \binom{n}{i} \cdot i \cdot \varphi^i$$
$$\Leftrightarrow \sum_{i=1}^{\frac{n-1}{2}} \binom{n}{i} \cdot \left(i\varphi^i - (n-i)\varphi^{n-i}\right) > n\varphi^n.$$

A sufficient condition for this to hold is that

$$i\varphi^i - (n-i)\varphi^{n-i} > 0, \forall i \in [(n-1)/2]$$

and

$$\frac{n-1}{2}\varphi^{\frac{n-1}{2}} - \frac{n+1}{2}\varphi^{\frac{n+1}{2}} \geqslant n\varphi^n.$$

For the former condition, it is sufficient to have

$$\frac{n-1}{2}\varphi^{\frac{n-1}{2}} - \frac{n+1}{2}\varphi^{\frac{n+1}{2}} > 0$$
$$\Leftrightarrow \varphi < \frac{(n-1)/2}{(n+1)/2} = 1 - \frac{2}{n+1}.$$

This is true because $\varphi < 1 - 4\ln n/(n+1)$. We can simplify the latter condition to

$$\frac{n-1}{2} \geqslant \frac{n+1}{2}\varphi + n\varphi^{\frac{n+1}{2}}.$$

For $\varphi \leqslant 1 - \frac{4\ln n}{n+1}$, we have

$$\begin{aligned}\frac{n+1}{2}\varphi + n\varphi^{\frac{n+1}{2}} &\leqslant \frac{n+1-4\ln n}{2} + ne^{-\frac{4\ln n}{n+1}\cdot\frac{n+1}{2}} \\ &= \frac{n+1-4\ln n}{2} + 1 \leqslant \frac{n-1}{2},\end{aligned}$$

where the first transition holds because $1 - x \leqslant e^{-x}$ for all $x$, and the last transition holds when $\ln n \geqslant 1$, which is true for $n \geqslant 3$.

**Case 2:** $k = (n-1)/2$**.** In this case, we want to show $\varphi_k^* \leqslant 2/(n+1)$. It is sufficient to show that for all $\varphi > 2/(n+1)$,

$$\sum_{i=k}^{\frac{n-1}{2}} \binom{i}{k}\binom{n}{i}\varphi^i < \sum_{i=\frac{n+1}{2}}^{n} \binom{i}{k}\binom{n}{i}\varphi^i.$$

For $k = (n-1)/2$, this reduces to

$$\binom{n}{\frac{n-1}{2}}\varphi^{\frac{n-1}{2}} < \sum_{i=\frac{n+1}{2}}^{n} \binom{i}{\frac{n-1}{2}}\binom{n}{i}\varphi^i.$$

It suffices to show that

$$\binom{n}{\frac{n-1}{2}}\varphi^{\frac{n-1}{2}} < \binom{\frac{n+1}{2}}{\frac{n-1}{2}}\binom{n}{\frac{n+1}{2}}\varphi^{\frac{n+1}{2}}.$$

Upon simplification, it is easy to see that this is true when $\varphi > \frac{2}{n+1}$.

**Case 3:** $k = (n-1)/4$**.** In this case, we want to show that $\varphi_k^* \in [1/4, 2/3]$. It is sufficient to show that

$$\sum_{i=\frac{n-1}{4}}^{\frac{n-1}{2}} \binom{i}{\frac{n-1}{4}}\binom{n}{i}\varphi^i > \sum_{i=\frac{n+1}{2}}^{n} \binom{i}{\frac{n-1}{4}}\binom{n}{i}\varphi^i,\ \forall\varphi < 1/4, \tag{6.7}$$

and

$$\sum_{i=\frac{n-1}{4}}^{\frac{n-1}{2}} \binom{i}{\frac{n-1}{4}}\binom{n}{i}\varphi^i < \sum_{i=\frac{n+1}{2}}^{n} \binom{i}{\frac{n-1}{4}}\binom{n}{i}\varphi^i,\ \forall\varphi > 2/3. \tag{6.8}$$

We show each inequality separately.

**Case 3a:** $\varphi < 1/4$**.** For Equation (6.7), it suffices to show that for all $\varphi < 1/4$ and $i \in \left\{0, 1, \ldots, \frac{n-1}{4}\right\}$, we have

$$\begin{aligned}
&\binom{\frac{n-1}{2} - i}{\frac{n-1}{4}} \binom{n}{\frac{n-1}{2} - i} \varphi^{\frac{n-1}{2} - i} > \\
&\binom{\frac{n+1}{2} + 2i}{\frac{n-1}{4}} \binom{n}{\frac{n+1}{2} + 2i} \varphi^{\frac{n+1}{2} + 2i} \\
&+ \binom{\frac{n+1}{2} + 2i + 1}{\frac{n-1}{4}} \binom{n}{\frac{n+1}{2} + 2i + 1} \varphi^{\frac{n+1}{2} + 2i + 1} \\
&\Leftarrow 1 > \frac{(\frac{n-1}{2} + i + 1) \cdots (\frac{n-1}{2} - 2i + 1)}{(\frac{n-1}{4} + 2i + 1) \cdots (\frac{n-1}{4} - i + 1)} \left(\frac{1}{4}\right)^{3i+1} \\
&+ \frac{(\frac{n-1}{2} + i + 1) \cdots (\frac{n-1}{2} - 2i)}{(\frac{n-1}{4} + 2i + 2) \cdots (\frac{n-1}{4} - i + 1)} \left(\frac{1}{4}\right)^{3i+2}
\end{aligned}$$

We show that both terms of the right hand side are less than $1/2$ and the inequality follows. First, for $i \in \left\{0, \ldots, \frac{n-1}{4}\right\}$, we have

$$\left(\frac{n-1}{4} + 2i + 1\right) \cdot \frac{1}{2} > \left(\frac{n-1}{2} + i + 1\right) \cdot \frac{1}{4},$$

and for $i \in \left\{0, \ldots, \frac{n-1}{4}\right\}$ and $j \in [3i]$, we have

$$\frac{n-1}{4} + 2i + 1 - j > \left(\frac{n-1}{2} + i + 1 - j\right) \cdot \frac{1}{4}.$$

This shows that the first term is less than $1/2$. To show that the second term is less than $1/2$, it is sufficient to notice that for $i \in \left\{0, \ldots, \frac{n-1}{4}\right\}$,

$$\left(\frac{n-1}{4} + 2i + 2\right) \cdot \frac{1}{2} > \left(\frac{n-1}{2} + i + 1\right) \cdot \frac{1}{4},$$

and for $i \in \left\{0, \ldots, \frac{n-1}{4}\right\}$ and $j \in [3i + 1]$,

$$\frac{n-1}{4} + 2i + 2 - j > \left(\frac{n-1}{2} + i + 1 - j\right) \cdot \frac{1}{4}.$$

**Case 3b:** $\varphi > 2/3$**.** We now show Equation (6.8) for $\varphi = 2/3$. Then, the equation clearly holds for $\varphi > 2/3$. First, note that for $i \in \left\{0, \ldots, \frac{n-1}{4} - 1\right\}$, it holds that

$$\begin{aligned}
&\binom{\frac{n+1}{2} + i}{\frac{n-1}{4}} \binom{n}{\frac{n+1}{2} + i} \left(\frac{2}{3}\right)^{\frac{n+1}{2} + i} \geqslant \\
&\binom{\frac{n-1}{2} - i}{\frac{n-1}{4}} \binom{n}{\frac{n-1}{2} - i} \left(\frac{2}{3}\right)^{\frac{n-1}{2} - i} \\
&\Leftrightarrow \left(\frac{2}{3}\right)^{2i+1} \geqslant \frac{\left(\frac{n-1}{4} + i + 1\right) \ldots \left(\frac{n-1}{4} - i + 1\right)}{\left(\frac{n-1}{2} + i + 1\right) \ldots \left(\frac{n-1}{2} - i + 1\right)}.
\end{aligned}$$

This is because for each $j \in \{0, \ldots, 2i\}$, we have

$$\frac{2}{3} \cdot \left( \frac{n-1}{2} + i + 1 - j \right) \geqslant \frac{n-1}{4} + i + 1 - j.$$

Next, we argue that

$$\begin{aligned}
&\binom{\frac{3n+1}{4}}{\frac{n-1}{4}} \binom{n}{\frac{3n+1}{4}} \left(\frac{2}{3}\right)^{\frac{3n+1}{4}} \\
&+ \binom{\frac{3n+1}{4}+1}{\frac{n-1}{4}} \binom{n}{\frac{3n+1}{4}+1} \left(\frac{2}{3}\right)^{\frac{3n+1}{4}+1} > \binom{n}{\frac{n-1}{4}} \left(\frac{2}{3}\right)^{\frac{n-1}{4}} \\
&\Leftrightarrow \frac{\left(\frac{3n+1}{4}\right) \ldots \left(\frac{n+1}{2}+1\right)}{(\frac{n-1}{4})!} (\frac{2}{3})^{\frac{n+1}{2}} \\
&\quad + \frac{\left(\frac{3n+1}{4}\right) \ldots \left(\frac{n+1}{2}+2\right)}{(\frac{n-1}{4}-1)!} (\frac{2}{3})^{\frac{n+1}{2}+1} > 1.
\end{aligned}$$

To prove this, we show that the first term on the LHS is greater than $7/8$, and the second term on the LHS is greater than $1/8$. First, we notice that

$$\left(\frac{2}{3}\right)^3 \cdot \frac{3n+1}{4} > \frac{7}{8} \cdot \left(\frac{n-1}{4}\right).$$

Next, for each $j \in [\frac{n-1}{4} - 1]$,

$$\left(\frac{2}{3}\right)^2 \cdot \left(\frac{3n+1}{4} - j\right) > \frac{n-1}{4} - j.$$

This is sufficient to show that the first term is greater than $7/8$. For the second term, note that

$$\left(\frac{2}{3}\right)^6 \cdot \frac{3n+1}{4} > \frac{1}{8} \cdot \left(\frac{n-1}{4} - 1\right),$$

and for each $j \in [\frac{n-1}{4} - 2]$,

$$\left(\frac{2}{3}\right)^2 \left(\frac{3n+1}{4} - j\right) > \frac{n-1}{4} - 1 - j.$$

Thus, the second term is greater than $1/8$. This concludes the proof. □

Let us consider the implications of Theorem 31 as $n \to \infty$. The first part implies that if we observe only a single $b \succ a$ sample, we should predict $a \succ b$ for any $\varphi < 1$. This makes sense because the $n-1$ unobserved votes vastly overshadow the single observed vote, and the prior places at least somewhat more probability on $a \succ b$ than on $b \succ a$.

The second part implies that if we observe $(n-1)/2$ votes (just a little less than a majority), we should predict $b \succ a$ for any $\varphi > 0$. This again makes sense because the probability that there is at least one $b \succ a$ in the remaining $(n+1)/2$ votes — sufficient to make $b \succ a$ the plurality outcome on the original profile — approaches 1.

The final part shows that the transition between $\varphi_k^* \approx 0$ and $\varphi_k^* \approx 1$ is not sudden; for $k = (n-1)/4$, the transition happens at $\varphi_k^*$ that is not arbitrarily close to either endpoint when $n$ is large.

| | Borda | Bucklin | Copeland | Harmonic | Maximin | Plurality | STV | Veto | Borda | Bucklin | Copeland | Harmonic | Maximin | Plurality | STV | Veto |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Borda | 1.40 | 1.26 | 0.43 | 1.65 | 0.59 | 0.61 | 0.79 | 0.59 | 8.32 | 5.73 | 3.83 | 6.49 | 4.93 | 3.12 | 4.03 | 3.32 |
| Bucklin | 1.27 | 1.39 | 0.38 | 1.37 | 0.53 | 0.46 | 0.65 | 0.54 | 5.67 | 8.16 | 2.67 | 3.54 | 3.55 | 1.59 | 2.24 | 2.54 |
| Copeland | 4.18 | 3.81 | 2.46 | 4.48 | 2.64 | 2.61 | 3.01 | 2.55 | 15.67 | 11.75 | 12.10 | 12.99 | 10.40 | 7.81 | 10.44 | 8.16 |
| Harmonic | 1.22 | 1.04 | 0.36 | 1.67 | 0.52 | 0.67 | 0.79 | 0.44 | 5.90 | 3.02 | 2.72 | 8.43 | 3.80 | 5.70 | 5.46 | 1.60 |
| Maximin | 1.59 | 1.47 | 0.51 | 1.82 | 0.77 | 0.73 | 0.90 | 0.71 | 6.77 | 5.05 | 3.03 | 5.73 | 6.59 | 3.13 | 4.17 | 3.12 |
| Plurality | 1.30 | 1.05 | 0.43 | 1.88 | 0.61 | 0.89 | 0.98 | 0.51 | 4.08 | 1.97 | 2.08 | 7.99 | 2.93 | 8.12 | 5.76 | 1.20 |
| STV | 1.36 | 1.16 | 0.49 | 1.86 | 0.62 | 0.84 | 1.19 | 0.54 | 5.01 | 2.76 | 3.13 | 7.21 | 3.77 | 5.46 | 10.66 | 1.54 |
| Veto | 1.33 | 1.23 | 0.45 | 1.35 | 0.59 | 0.52 | 0.66 | 0.89 | 4.48 | 3.63 | 2.19 | 2.56 | 3.01 | 1.31 | 1.78 | 8.34 |
| Borda | 99.84 | 99.80 | 99.95 | 98.69 | 62.54 | 40.35 | 40.65 | 40.43 | 100 | 100 | 100 | 100 | 99.87 | 97.35 | 97.35 | 97.60 |
| Bucklin | 99.84 | 99.80 | 99.95 | 98.69 | 62.54 | 40.35 | 40.65 | 40.43 | 100 | 100 | 100 | 100 | 99.87 | 97.35 | 97.35 | 97.60 |
| Copeland | 99.84 | 99.80 | 99.95 | 98.69 | 62.54 | 40.35 | 40.65 | 40.43 | 100 | 100 | 100 | 100 | 99.87 | 97.35 | 99.35 | 97.60 |
| Harmonic | 99.84 | 99.80 | 99.95 | 98.69 | 62.54 | 40.35 | 40.65 | 40.43 | 100 | 100 | 100 | 100 | 99.87 | 97.35 | 97.35 | 97.60 |
| Maximin | 99.84 | 99.80 | 99.95 | 98.69 | 62.54 | 40.35 | 40.65 | 40.43 | 100 | 100 | 100 | 100 | 99.87 | 97.35 | 97.35 | 97.60 |
| Plurality | 99.64 | 99.60 | 99.75 | 98.52 | 62.50 | 40.42 | 40.72 | 40.34 | 99.9 | 99.9 | 99.9 | 99.9 | 99.77 | 97.66 | 97.66 | 97.50 |
| STV | 99.64 | 99.60 | 99.75 | 98.52 | 62.50 | 40.42 | 40.72 | 40.34 | 99.9 | 99.9 | 99.9 | 99.9 | 99.77 | 97.66 | 97.66 | 97.50 |
| Veto | 99.84 | 99.80 | 99.95 | 98.69 | 62.54 | 40.35 | 40.65 | 40.55 | 99.9 | 99.9 | 99.9 | 99.9 | 99.77 | 97.25 | 97.25 | 97.75 |

Figure 6.1: Average-case predictability of different voting rules $f_1$ (rows) using different voting rules $f_2$ (columns) under the uniform distribution (top) and the Mallows model with $\varphi = 1/3$ (bottom) with $k = 50$ (left) and $k = 500$ (right).

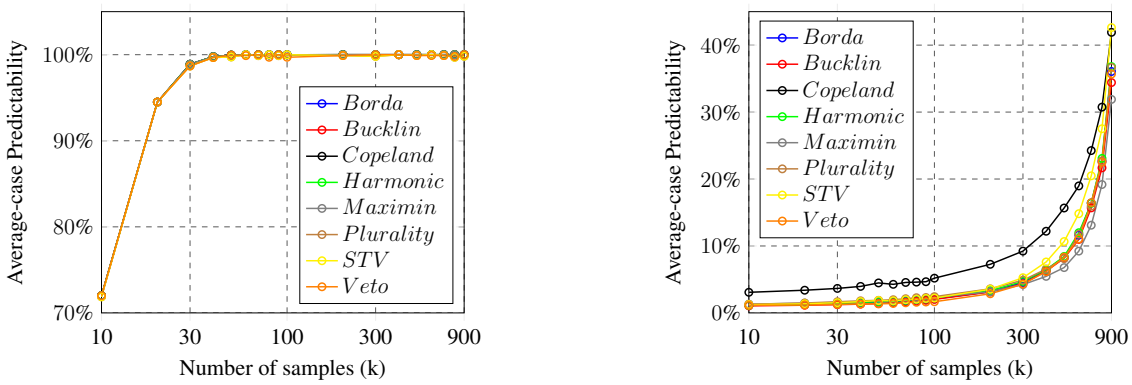


Figure 6.2: Average-case predictability of different voting rules as a function of the number of samples $k$, under the Mallows model with $\varphi = 1/3$ (left) and under the uniform distribution (right).

## 6.5 Experiments

In this section, we conduct experiments to measure the predictability of popular voting rules in the average case.[3] We consider profiles $\vec{\sigma}_n$ with $n = 1,000$ voters and $m = 5$ alternatives. We use two distributions to draw i.i.d. rankings in $\vec{\sigma}_n$: the Mallows model with $\varphi = 1/3$ and the Mallows model with $\varphi = 1$. The latter case corresponds to the uniform distribution, i.e. the samples are not concentrated around the ground truth, and in contrast, the former case corresponds to samples that are

[3]Refer to the book by Brandt et al. [2016] for definitions.

quite concentrated around the ground truth. We refer to the former distribution as "Mallows distribution" and to the latter distribution as "uniform distribution". We average our results across $10^6$ draws of profile $\vec{\sigma}_n$.

Figure 6.1 shows the average predictability of different voting rules $f_1$ (rows) using different voting rules $f_2$ (columns), under the uniform distribution (tables on the top) and under the Mallows distribution (tables on the bottom), with $k = 50$ (tables on the left) and with $k = 500$ (tables on the right).[4] The entries in the table indicate the percentage of instances on which prediction was successful. Generally, we observe that prediction accuracy increases as the prior becomes more concentrated and as the number of samples $k$ increases, as expected. We also note a few peculiarities. Under the uniform distribution with $k = 50$, the harmonic rule is the best predictor of every voting rule (except Bucklin), although the prediction accuracy is small. As $k$ increases to $500$, however, each voting rule (except Copeland and maximin) becomes the best predictor of itself. Under the Mallows distribution, it is evident that Borda, Bucklin, Copeland, and the harmonic rule predict other voting rules well — often because they return fewer ties — while maximin, plurality, STV, and veto perform worse.

Figure 6.2 shows the average-case predictability of different voting rules $f_1$ (using the best voting rule $f_2$ from the same list) as a function of the number of samples, under the Mallows distribution (left) and under the uniform distribution (right). Once again, more concentrated prior and more samples allow greater predictability. The effect of the prior is significant: under the Mallows distribution, observing just $3\%$ of the votes allows predicting every voting rule with at least $98\%$ accuracy, while the same number of samples under the uniform distribution does not allow predicting any voting rule with more than $4\%$ accuracy.

## 6.6 Winner Prediction

While Theorems 26 and 27 paint an extremely pessimistic picture of predictability of positional scoring rules, this could be because we want to predict the entire ranking of alternatives returned by the rule. But sometimes we may be interested in predicting only the top alternative. Could this lead to more optimistic results? The answer is yes and no. We show that the worst-case accuracy of predicting the Borda count winner using Borda count itself is still zero.

**Theorem 32.** *Let $n \geqslant 2$, $m \geqslant 4$, and $k \in [n-1]$ such that $n$ and $k$ have different parity. Then, $\mathrm{acc}(f_{Borda}, f_{Borda}) = 0$, where $f_{Borda}$ denotes Borda count.*

*Proof.* We consider cases of even and odd $k$ and for each case we construct a profile on which Borda count cannot predict its own winner.

**Odd $n$, even $k$:** We start with the case where $n$ is odd and $k$ is even. Consider the following profile $\vec{\sigma}_n$. Each row represents a ranking where alternatives are listed from left to right in the most preferred to least preferred order. The first two rankings appear $(n-1)/2$ times each, and the third ranking appears once. Alternatives not shown appear in an arbitrary order in the bottom.

We denote with $\sigma_1$, $\sigma_2$ and $\sigma_3$ the ranking of the first, second and third rows, respectively. Under Borda rule, clearly winner is $a_3$.

[4] For the Bucklin rule, we define the Bucklin score of an alternative as the smallest $t$ such that a majority of voters rank the alternative in the first $t$ positions. Alternatives are first compared by their Bucklin score (lower is better), and alternatives with the same Bucklin score $t$ are compared by the number of voters who rank them in the first $t$ positions (higher is better).

| | |
|---|---|
| $\frac{n-1}{2}$ votes | $a_1 \succ a_3 \succ a_2 \succ a_4 \succ \ldots$ |
| $\frac{n-1}{2}$ votes | $a_4 \succ a_2 \succ a_3 \succ a_1 \succ \ldots$ |
| 1 vote | $a_3 \succ a_1 \succ a_4 \succ a_2 \succ \ldots$ |

We show that for every sample $\vec{\pi}_k \subset \vec{\sigma}_n$ of size $k$, $f_{\text{Borda}}(\vec{\pi}_k) \neq \{a_3\}$. Let $x_1$, $x_2$, and $x_3$ denote the number of times $\sigma_1$, $\sigma_2$, and $\sigma_3$ appear in $\vec{\pi}_k$, respectively. Note that $x_3 \in \{0, 1\}$.

In order for $a_3$ to be the unique winner under $f_{\text{Borda}}$ on $\vec{\pi}_k$, it must defeat $a_1$. Hence, we need

$$\begin{aligned} &x_1 \cdot (m-2) + x_2 \cdot (m-3) + x_3 \cdot (m-1) \\ &\quad > x_1 \cdot (m-1) + x_2 \cdot (m-4) + x_3 \cdot (m-2) \\ &\Rightarrow x_2 + x_3 > x_1. \end{aligned}$$

Similarly, $a_3$ must also defeat $a_4$. Hence, we also need

$$\begin{aligned} &x_1 \cdot (m-2) + x_2 \cdot (m-3) + x_3 \cdot (m-1) \\ &\quad > x_1 \cdot (m-4) + x_2 \cdot (m-1) + x_3 \cdot (m-3) \\ &\Rightarrow x_1 + x_3 > x_2. \end{aligned}$$

Given $x_1 + x_3 > x_2$ and $x_2 + x_3 > x_1$, we can derive $x_3 > 0$, i.e., $x_3 = 1$. But then, we have $x_1 + 1 > x_2$ and $x_2 + 1 > x_1$, which implies $x_1 = x_2$. In this case, $|\vec{\pi}_k| = x_1 + x_2 + 1$ is odd, which contradicts the fact that $k$ is even.

**Even $n$, odd $k$:** We now consider the case of even $n$ and odd $k$. Consider the following profile $\vec{\sigma}_n$.

| | |
|---|---|
| $\frac{n-1}{2}$ votes | $a_1 \succ a_3 \succ a_2 \succ a_4 \succ \ldots$ |
| $\frac{n-1}{2}$ votes | $a_4 \succ a_2 \succ a_3 \succ a_1 \succ \ldots$ |
| 1 vote | $a_3 \succ a_2 \succ a_1 \succ a_4 \succ \ldots$ |
| 1 vote | $a_4 \succ a_2 \succ a_1 \succ a_3 \succ \ldots$ |

We again denote with $\sigma_1$, $\sigma_2$, $\sigma_3$ and $\sigma_4$ the rankings in rows 1, 2, 3 and 4, respectively. When Borda count is applied on $\vec{\sigma}_n$, clearly $a_2$ is the unique winner.

We now argue that for every sample $\vec{\pi}_k \subset \vec{\sigma}_n$ of size $k$, $f_{\text{Borda}}(\vec{\pi}_k) \neq \{a_2\}$. Again, let $x_1$, $x_2$, $x_3$, and $x_4$ denote the number of times $\sigma_1$, $\sigma_2$, $\sigma_3$, and $\sigma_4$ appear in $\vec{\pi}_k$, respectively. Note that $x_3, x_4 \in \{0, 1\}$. In order for $a_2$ to be the unique winner under $f_{\text{Borda}}$ on $\vec{\pi}_k$, it must defeat $a_3$. Hence, we need

$$\begin{aligned} &x_1 \cdot (m-3) + x_2 \cdot (m-2) + x_3 \cdot (m-2) + x_4 \cdot (m-2) \\ &> x_1 \cdot (m-2) + x_2 \cdot (m-3) + x_3 \cdot (m-1) + x_4 \cdot (m-4) \\ &\Rightarrow x_2 + 2 \cdot x_4 > x_1 + x_3. \end{aligned} \tag{6.9}$$

Similarly, $a_2$ must also defeat $a_4$. Hence, we also need

$$\begin{aligned} &x_1 \cdot (m-3) + x_2 \cdot (m-2) + x_3 \cdot (m-2) + x_4 \cdot (m-2) \\ &> x_1 \cdot (m-4) + x_2 \cdot (m-1) + x_3 \cdot (m-4) + x_4 \cdot (m-1) \\ &\Rightarrow x_1 + 2 \cdot x_3 > x_2 + x_4. \end{aligned} \tag{6.10}$$

**Case 1:** $x_3 = x_4 = 0$**.** From Equation (6.9), we obtain $x_2 > x_1$, while from Equation (6.10), we obtain $x_1 > x_2$, which is clearly a contradiction.

**Case 2:** $x_3 = 1$ **and** $x_4 = 0$**.** From Equation (6.9), we obtain $x_2 > x_1 + 1$, while from Equation (6.10), we obtain $x_1 + 2 > x_2$, which is a contradiction.

**Case 3:** $x_3 = 0$ **and** $x_4 = 1$**.** This case leads to a contradiction in a manner similar to Case 2, so we omit the details.

**Case 4:** $x_3 = 1$ **and** $x_4 = 1$**.** From Equation (6.9) we obtain that $x_2 + 2 > x_1 + 1$, while from Equation (6.10) we obtain $x_1 + 2 > x_2 + 1$. This means that $x_1 = x_2$, which implies $|\vec{\pi}_k| = x_1 + x_2 + 2$ is even, which in turn contradicts the fact that $k$ is odd. □

But the case of the $r$-approval rule *is* more optimistic.

**Theorem 33.** *For $r \in [m-1]$ and $k \geqslant r$, we have $\mathrm{acc}(f_{r\text{-app}}, f_{r\text{-app}}) > 0$, where $f_{r\text{-app}}$ denotes the $r$-approval voting rule.*

*Proof.* Fix $r \in [m-1]$ and $k \geqslant r$. Consider any profile $\vec{\sigma}_n$. Let $W = f_{r\text{-app}}(\vec{\sigma}_n)$ denote the set of winners under $f_{r\text{-app}}$ on $\vec{\sigma}_n$, and let $\overline{W} = A \setminus W$. Let $x_i$ denote the $r$-approval score of $a_i$ in $\vec{\sigma}_n$ for each $i \in [m]$.

We now show that there exists a sample $\vec{\pi}_k \in \mathcal{S}_k(\vec{\sigma}_n)$ such that $f_{r\text{-app}}(\vec{\pi}_k) \subseteq f_{r\text{-app}}(\vec{\sigma}_n)$. Without loss of generality, assume that $a_1 \in W$. Hence, we have $x_1 > x_j$ for all $a_j \in \overline{W}$.

Let $\vec{\rho} \subseteq \vec{\sigma}_n$ denote the subset of votes in which $a_1$ is approved (i.e. it appears in the top $r$ positions). If $k \geqslant |\vec{\rho}|$, then we can simply choose an arbitrary $\vec{\pi}_k \in \mathcal{S}_k(\vec{\sigma}_n)$ such that $\vec{\rho} \subseteq \vec{\pi}_k$. In this case, the score of $a_1$ would be exactly equal to $x_1 = |\vec{\rho}|$, while the score of any $a_j \in \overline{W}$ will be at most $x_j < x_1$. This implies $f_{r\text{-app}}(\vec{\pi}_k) \subseteq W$, as desired.

Now, suppose $r \leqslant k \leqslant |\vec{\rho}|$. Note that the previous argument establishes that for $k = |\vec{\rho}|$, we can use $\vec{\pi}_k = \vec{\rho}$, and have $f_{r\text{-app}}(\vec{\pi}_k) \subseteq W$. For $k < |\vec{\rho}|$, we show that one can start from $\vec{\rho}$, and iteratively remove one vote at a time — until $k$ votes remain — such that the set of winners is always a subset of $W$.

More specifically, we show that given $k \geqslant r$, and any $\vec{\pi}_{k+1} \subseteq \vec{\rho}$ with $|\vec{\pi}_{k+1}| = k+1$ and $f_{r\text{-app}}(\vec{\pi}_{k+1}) \subseteq W$, there exists $\vec{\pi}_k \subset \vec{\pi}_{k+1}$ with $|\vec{\pi}_k| = k$ and $f_{r\text{-app}}(\vec{\pi}_k) \subseteq W$.

Suppose this is false for some $\vec{\pi}_{k+1}$. That is, removing any single ranking from $\vec{\pi}_{k+1}$ results in at least one alternative from $\overline{W}$ becoming a winner. For each $\sigma \in \vec{\pi}_{k+1}$, let us denote $\vec{\pi}_k^{\sigma} = \vec{\pi}_{k+1} \setminus \{\sigma\}$ and $S(\sigma) = f_{r\text{-app}}(\vec{\pi}_k^{\sigma}) \cap \overline{W}$. Then, our assumption implies that $S(\sigma) \neq \emptyset$ for each $\sigma \in \vec{\pi}_{k+1}$.

We further show that $S(\sigma) \cap S(\sigma') = \emptyset$ for distinct $\sigma, \sigma' \in \vec{\pi}_{k+1}$.[5] Note that for each $a_j \in S(\sigma)$, $a_j$ is a winner under $\vec{\pi}_k^{\sigma}$ but not under $\vec{\pi}_{k+1}$. Also, note that the $r$-approval score of $a_1$ under $\vec{\pi}_k^{\sigma}$ and $\vec{\pi}_{k+1}$ is $k$ and $k+1$, respectively, since $\vec{\pi}_{k+1} \subseteq \vec{\rho}$. Hence, it must be the case that $a_j$ is approved in every ranking in $\vec{\pi}_k^{\sigma}$ but not in $\sigma$. Hence, $a_j \notin S(\sigma')$ for every other ranking $\sigma'$ in $\vec{\pi}_{k+1}$.

Finally, let us partition $\overline{W}$ into two sets $S_{\text{win}}$ and $S_{\text{lose}}$ such that $S_{\text{win}} = \cup_{\sigma \in \vec{\pi}_{k+1}} S(\sigma)$ and $S_{\text{lose}} = \overline{W} \setminus S_{\text{win}}$. Let $t = \min_{\sigma \in \vec{\pi}_{k+1}} |S(\sigma)|$. Then, $|S_{\text{win}}| \geqslant t \cdot (k+1)$. Further, consider a ranking $\sigma \in \vec{\pi}_{k+1}$ under which exactly $t$ of the unapproved alternatives are in $S_{\text{win}}$. In this ranking, the remaining $m - r - t$ unapproved alternatives must be in $S_{\text{lose}}$. Hence, $|S_{\text{lose}}| \geqslant m - r - t$. Thus, we have

[5] Here, even two copies of the same ranking are referred to as distinct.

$\overline{W} \geqslant t \cdot (k+1) + m - r - t$. However, since $a_1 \in W$, we also have $\overline{W} \leqslant m - 1$. Combining the two inequalities, we get

$$\begin{aligned} t \cdot (k+1) + m - r - t \leqslant m - 1 &\Rightarrow t \cdot k \leqslant r - 1 \\ &\Rightarrow k \leqslant r - 1, \end{aligned}$$

where the last inequality holds because $t \geqslant 1$ (since $S(\sigma) \neq \emptyset$ for each $\sigma \in \vec{\pi}_{k+1}$). However, this is a contradiction since we assumed $k \geqslant r$. □

It is easy to notice that the lower bound of $r$ is tight. For example, consider the following preference profile with $m$ alternatives and $r$ votes.

| Approved | Not Approved |
|---|---|
| $a_1 \succ a_2 \succ a_3 \succ \ldots \succ a_r$ | $a_{r+1} \succ a_{r+2} \succ \ldots \succ a_m$ |
| $a_1 \succ a_{r+1} \succ a_2 \succ \ldots \succ a_{r-1}$ | $a_r \succ a_{r+2} \succ \ldots \succ a_m$ |
| $\vdots$ | $\vdots$ |
| $a_1 \succ a_3 \succ a_4 \succ \ldots \succ a_{r+1}$ | $a_2 \succ a_{r+2} \succ \ldots \succ a_m$ |

Basically, alternative $a_1$ is approved in each vote, and in each of the $r$ votes, a distinct subset of $r-1$ alternatives from the set $\{a_2, \ldots, a_{r+1}\}$ of $r$ alternatives is approved. It is easy to check that $a_1$ is the unique $r$-approval winner on this profile. However, there is no subset of $r-1$ (or fewer) votes under which $a_1$ is the unique $r$-approval winner.

## 6.7 Discussion

Predicting election outcomes using limited information is a broad research agenda, and specifically predicting the entire ranking of alternatives returned by the rule is required in several real-world applications, e.g., Borda count is used to rank college football teams in the Associated Press poll [Levin and Nalebuff, 1995] and to rank students in MOOCs [Caragiannis et al., 2019a]. While our work makes progress towards painting the full picture, there are a number of areas yet unexplored. The most immediate direction is to fill the gaps in our results, e.g., analyzing the accuracy of plurality and veto predicting themselves (Theorems 29 and 30) for all values of $k$, and extending the average-case analysis to heterogeneous samples, all values of $k$, and more than two alternatives. The next step would be to study other voting rules (e.g. Copeland's or Kemeny's method) and other models of sampling votes (e.g. when each voter $i$ independently participates in the poll with probability $p_i$). In Section 6.6, we argue that while predicting the winner under $k$-approval rules seems plausible, predicting the Borda winner is still difficult. Tracing the exact predictability of winner under positional scoring rules remains an interesting direction for future work.

Finally, we can also consider the use of limited information to make good collective decisions in other frameworks of voting. For example, in the implicit utilitarian voting framework [Procaccia and Rosenschein, 2006, Boutilier et al., 2015], where the goal is to find an alternative with small distortion, how small can we make the expected distortion given only sampled votes?

# Part III

# Algorithmic Fairness

# Chapter 7

# Proportionally Fair Clustering

## 7.1 Introduction

Machine learning algorithms are increasingly being used for decision making in applications where they affect human lives; popular examples include resume screening, evaluation of loan applications, bail decisions, etc. [Parkes and Vohra, 2019]. This has led to growing concern as to whether these algorithms, which may view humans as "data points", treat them fairly [Angwin et al., 2016, O'Neil, 2016]. Consequently, research on designing fair machine learning algorithms is proliferating [Narayanan, 2018, Mehrabi et al., 2021].

Much of this literature focuses on fairness in classification [Zemel et al., 2013, Zafar et al., 2019, Hossain et al., 2020], but the study of fairness in other settings such as regression [Agarwal et al., 2019] and clustering [Chen et al., 2019] is also on the rise. In this chapter, we focus on fairness in clustering, specifically, in centroid clustering. In this problem, we are given a set $\mathcal{N}$ of data points in a metric space, and a set $\mathcal{M}$ of possible locations for cluster centers in the same metric space. Given $k \in \mathbb{N}$, the task is to select a set $X \subseteq \mathcal{M}$ consisting of $|X| = k$ cluster centers and assign each data point to a cluster center — usually the closest — with the goal that data points are close to their assigned cluster centers. Clustering has diverse applications in market research, pattern recognition, data analysis, image segmentation, and facility location. In applications like image segmentation or market research, often the goal is to simply identify different clusters of points among the data. However, in facility location [Fotakis, 2006, Liao and Guo, 2008, Gupta et al., 2016], where data points may represent locations of houses in the neighborhood and cluster centers may represent locations where public facilities (such as parks) will be built, it is of paramount importance that the facilities be distributed to fairly serve the population.

Adapting an example given by Chen et al. [2019], imagine that there is a dense urban area with a population of $10{,}000$, and far from it, there are $10$ small communities with a population of $100$ each. The communities are closer to each other compared to how far they are from the urban area, but still well distinguished. With $k = 11$, a standard clustering algorithm such as $k$-means would identify the urban area as one cluster, and each small community as one cluster. However, building just one park that serves $10{,}000$ people in the urban area, while each community of $100$ people gets its own park violates the principle of *equal entitlement* [Moulin, 2003]; this principle would suggest that when allocating $11$ parks among a total of $11{,}000$ people, the urban area consisting of $10{,}000$ people should be allocated their proportional share of $10$ parks, and one park should serve the $10$ smaller communities consisting of

1,000 people altogether.

This notion of what a group deserves — *group fairness* — has been extensively studied in machine learning, and a variety of definitions have been proposed [Calders et al., 2009, Hardt et al., 2016, Narayanan, 2018, Hossain et al., 2020, Ustun et al., 2019]. Borrowing from a long line of literature on fair resource allocation [Varian, 1974, Shapley and Scarf, 1974, Fain et al., 2018, Conitzer et al., 2019], Chen et al. [2019] proposed a novel definition of fairness in clustering that perfectly fits our motivation. Given a metric $d$ over a set $\mathcal{N}$ of $n$ points and a set $\mathcal{M}$ of feasible cluster centers, they say that a $k$-clustering $X$ satisfies *proportional fairness* if there is no group of points $S \subseteq \mathcal{N}$ with $|S| \geqslant n/k$ and a new cluster location $y \in \mathcal{M}$ such that $d(i, y) < \min_{x \in X} d(i, x)$ for each member $i \in S$. Most fairness definitions in machine learning protect groups of individuals that are pre-defined based on certain *protected attributes* [Backurs et al., 2019, Bera et al., 2019, Chierichetti et al., 2017, Rösner and Schmidt, 2018] or that are sufficiently large [Kearns et al., 2018]. In contrast, proportional fairness guarantees fairness to arbitrarily defined groups of all possible sizes. This may be helpful given recent observations that protecting groups defined based on individual attributes may allow an algorithm to circumvent fairness [Kearns et al., 2018], or that information about which groups to protect may not be known in advance [Hashimoto et al., 2018]. For references to other related work, we direct the reader to the work of Chen et al. [2019].

## Our Contribution

We build upon the work of Chen et al. [2019]. While their work considers the metric $d$ and the set of feasible cluster center locations $\mathcal{M}$ to be arbitrary (and $|\mathcal{M}|$ to be typically finite), we focus on the case where the metric consists of usual distance functions such as $L^1$, $L^2$, or $L^\infty$ over $\mathbb{R}^t$, and cluster centers can be placed anywhere in the infinite metric space (i.e. $\mathcal{M} = \mathbb{R}^t$). While this change is seemingly simple, the infinite cardinality of $\mathcal{M}$ requires new algorithmic tools and generalization bounds, which we provide in this work. In some cases, we show that this in fact allows us to provide stronger approximation guarantees.

In Section 7.3, we analyze the greedy capture algorithm introduced by Chen et al. in the case where $d \in \{L^1, L^2, L^\infty\}$ and $\mathcal{M} = \mathbb{R}^t$. Chen et al. show that the algorithm provides $1 + \sqrt{2} \approx 2.414$ approximation to proportional fairness for all metric spaces. We show that for $d = L^2$ and $\mathcal{M} = \mathbb{R}^t$, it actually provides a better 2-approximation. We prove this via a refinement of the result of Chen et al.: we express the approximation ratio obtained by the algorithm in terms of a new characteristic of the metric that we term *Apollonius radius*, and show that this radius is small for the $L^2$ distance, allowing us to achieve a better approximation ratio. However, we show that for $L^1$ and $L^\infty$, the approximation ratio of greedy capture is no better than $1 + \sqrt{2}$.

In Section 7.4, we provide universal lower bounds which apply to all algorithms. Specifically, we show that for $d = L^2$ and $\mathcal{M} = \mathbb{R}^t$, no algorithm achieves better than $2/\sqrt{3} \approx 1.155$ approximation ratio, whereas for $d \in \{L^1, L^\infty\}$ and $\mathcal{M} = \mathbb{R}^t$, we get a lower bound of 1.4. Note that these lower bounds are existential, and not computational.

In Section 7.5, we consider the case where $\mathcal{M}$ is the set of nodes of an unweighted graph, and $d$ measures the shortest distance between two nodes on the graph. When the graph is a tree, we show that an exact proportionally fair clustering necessarily exists, and provide an efficient algorithm to find one. When the graph is arbitrary, but $k \geqslant n/2$ clusters need to be placed, we show that a proportionally fair clustering again necessarily exists and can be computed efficiently. Whether an exact proportionally fair

clustering exists for all graphs remains an interesting open question.

Next, in Section 7.6, we show that for $d = L^2$ and $\mathcal{M} = \mathbb{R}^t$ for $t \geqslant 2$, checking whether a proportional clustering exists is NP-hard. When $t$ is large, even implementing the greedy capture algorithm is NP-hard. However, this problem becomes efficiently solvable when $t$ is constant, and when $t$ is large, using a PTAS for an important sub-routine of greedy capture, we can efficiently compute a $2 \cdot (1 + \epsilon)$-proportionally fair clustering for any fixed $\epsilon > 0$.

Finally, in Section 7.7, we consider the problem of generalization: would a clustering that is proportionally fair with respect to samples drawn from $\mathcal{N}$ remain (approximately) proportionally fair with respect to the entire set $\mathcal{N}$? Chen et al. provide a positive answer for the case when $\mathcal{M}$ is finite. Using the framework of VC dimension, we show that, when $d = L^2$, the answer remains positive even when $\mathcal{M} = \mathbb{R}^t$.

## 7.2 Preliminaries

Let $\mathcal{N}$ be a set of $n$ data points (or agents), which lie in a metric space $(\mathcal{X}, d)$, where $d : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ is a distance function satisfying the triangle inequality. For most of this work, we consider the case where $\mathcal{X} = \mathbb{R}^t$ for some $t \in \mathbb{N}$, but in Section 7.5, we consider the case where $\mathcal{X}$ is the set of nodes of a graph. We also focus on special distance functions such as the Euclidean distance ($L^2$), the Manhattan distance ($L^1$), and the $L^\infty$ distance. Let $\mathcal{M} \subseteq \mathcal{X}$ be a the set of locations where cluster centers can be placed. In this work, we focus on the case where $\mathcal{M} = \mathcal{X}$. Let $k \in \mathbb{N}$. We denote by $[\mathcal{M}]^k$ the set of all subsets of $\mathcal{M}$ of size $k$. A $k$-clustering is a set $X \in [\mathcal{M}]^k$ containing $k$ locations for cluster centers. We refer to each $x \in X$ as an *open* cluster center.

The cost to agent $i \in \mathcal{N}$ induced by a cluster center $x \in \mathcal{M}$ is the distance $d(i, x)$, and the cost to agent $i \in \mathcal{N}$ induced by a $k$-clustering $X$ is the minimum distance from $i$ to any cluster center, i.e., $d(i, X) \triangleq \min_{x \in X} d(i, x)$. Agent $i$ is interested in minimizing her cost.

A set of points $S \subseteq \mathcal{N}$ containing at least $\lceil n/k \rceil$ is entitled at least one cluster center. If they can find a new cluster center that is better for each of them, we consider it a violation of fairness. Formally:

**Definition 13.** *Given $\rho \geqslant 1$, we say a $k$-clustering $X \in [\mathcal{M}]^k$ is* $\rho$-proportionally fair *if there is no $S \subseteq \mathcal{N}$ with $|S| \geqslant \lceil n/k \rceil$ and $y \in \mathcal{M}$ such that $\rho \cdot d(i, y) < d(i, X)$ for all $i \in S$. If such a coalition $S$ and point $y$ exist, we refer to $S$ as a* blocking coalition *and $y$ as the center that they* deviate *to. When $\rho = 1$, we simply call this proportional fairness.*

The reason to define $\rho$-proportionally fair solutions for $\rho > 1$ is because (exactly) proportionally fair solutions may not exist [Chen et al., 2019].

In Section 7.5, we consider the problem of proportional fairness on graphs. Specifically, given an undirected graph $G = (V, E)$, we assume that $\mathcal{N} \subseteq V$, $\mathcal{M} = V$ (i.e. every node in the graph is a feasible cluster center), and the distance between two nodes $u, v \in V$, denoted by $d(u, v)$, is the length of the shortest path connecting them. Notice that $d$ satisfies the triangle inequality.

## 7.3 Greedy Capture

In this section, we study the greedy capture algorithm defined by Chen et al. [2019]. Put succinctly, the algorithm starts with $X = \emptyset$. It grows a ball at every location in $\mathcal{M}$ at the same rate. As soon as a

ball contains at least $\lceil n/k \rceil$ points, the corresponding center is added to $X$ and all the points covered by the ball are removed. As balls continue growing, balls centered at previously added locations in $X$ also continue growing with them and any new points covered by such balls are immediately removed. We refer the reader to the work of Chen et al. for full description of the algorithm. They show that for any metric space, greedy capture is guaranteed to find a $(1+\sqrt{2})$-proportionally fair clustering.

We begin by providing a refined analysis of greedy capture by expressing the approximation ratio in terms of a characteristic of the metric we call the Apollonius radius.

**Definition 14.** *Given $\rho \geqslant 1$, the $\rho$-Apollonius radius of a metric $(\mathcal{X}, d)$ is defined as*

$$A_{\mathcal{X},d}(\rho) = \sup_{x,y \in \mathcal{X}} \Delta(\rho, x, y)/d(x, y),$$

*where $\Delta(\rho, x, y)$ is the radius of the smallest ball centered at some point in $\mathcal{X}$ that contains the entire set $\{p \in \mathcal{X} : \rho \cdot d(p, y) \leqslant d(p, x)\}$.*

The reason that we term it the Apollonius radius is because the renowned Greek geometer Apollonius of Perga was famously interested in the set $\{p \in \mathbb{R}^t : \rho \cdot d(p, y) \leqslant d(p, x)\}$ for $d = L^2$, and showed that this set is a ball already when $\rho > 1$. This special structure of $L^2$ is what will allow us to achieve a better approximation guarantee for it.

**Theorem 34.** *For any metric $(\mathcal{X}, d)$ and $\mathcal{M} = \mathcal{X}$, greedy capture finds a $\rho$-proportionally fair clustering, where $\rho \geqslant 1$ is the smallest positive number satisfying $A_{\mathcal{X},d}(\rho) \cdot \frac{\rho+1}{\rho} \leqslant 1$.*

*Proof.* Let $X$ be the clustering returned by the algorithm. Suppose that $X$ is not $\rho$-proportionally fair for some $\rho$. Then, there exist $S \subseteq \mathcal{N}$ with $|S| \geqslant \lceil n/k \rceil$ and $y \in \mathcal{M}$ such that $\rho \cdot d(j, y) < d(j, X)$ for all $j \in S$.

Note that this implies $y \notin X$. Let $i$ be the first point in $S$ that was covered during the execution of greedy capture; suppose it was covered by a ball located at $x \in X$.

Note that for each $j \in S$, we have $\rho \cdot d(j, y) < d(j, X) \leqslant d(j, x)$. Hence, we get that

$$S \subseteq \{p \in \mathcal{X} : \rho \cdot d(p, y) < d(p, x)\}.$$

Therefore, by definition of $\rho$-Apollonius radius, there exists a ball of radius at most $A_{\mathcal{X},d}(\rho) \cdot d(x, y)$ that contains all points in $S$. Since $i$ was the first point in $S$ covered by the ball centered at $x$, we must have $d(i, x) \leqslant A_{\mathcal{X},d}(\rho) \cdot d(x, y)$, otherwise greedy capture would have selected the ball of radius at most $A_{\mathcal{X},d}(\rho) \cdot d(x, y)$ covering all points of $S$ prior to $i$ being covered by the ball centered at $x$. Further, since $i \in S$, we also have $d(i, y) < d(i, x)/\rho$. Using the triangle inequality, we get

$$d(x, y) \leqslant d(i, x) + d(i, y) < d(i, x) \cdot \frac{\rho+1}{\rho} \leqslant A_{\mathcal{X},d}(\rho) \cdot \frac{\rho+1}{\rho} \cdot d(x, y) \Rightarrow A_{\mathcal{X},d}(\rho) \cdot \frac{\rho+1}{\rho} > 1.$$

We have proved that if $X$ is not $\rho$-proportionally fair, then $A_{\mathcal{X},d}(\rho) \cdot \frac{\rho+1}{\rho} > 1$. Hence, whenever $A_{\mathcal{X},d}(\rho) \cdot \frac{\rho+1}{\rho} \leqslant 1$, we have that $X$ is $\rho$-proportionally fair. □

Let us argue that Theorem 34 is in fact a refinement of the $(1+\sqrt{2})$-approximation proved by Chen et al. [2019] that holds for all metrics.

**Theorem 35.** *For any metric $(\mathcal{X}, d)$, the $\rho$-Apollonius radius is $A_{\mathcal{X},d}(\rho) \leqslant \frac{1}{\rho-1}$. Hence, greedy capture finds a $(1+\sqrt{2})$-proportionally fair clustering for every metric.*

*Proof.* Consider the set $\{p \in \mathcal{X} : \rho \cdot d(p,y) \leqslant d(p,x)\}$. For any point $p$ in this set, we have that $\rho \cdot d(p,y) \leqslant d(p,x) \leqslant d(p,y) + d(x,y)$ by the triangle inequality. Hence, $d(p,y) \leqslant d(x,y) \cdot \frac{1}{\rho-1}$ for all $p$ in the set. Thus, a ball centered at $y \in \mathcal{M}$ with radius $\frac{d(x,y)}{\rho-1}$ certainly covers the entire set. Hence, $A_{\mathcal{X},d}(\rho) \leqslant \frac{1}{\rho-1}$. Next, for $\rho = 1+\sqrt{2}$, we have

$$A_{\mathcal{X},d}(\rho) \cdot \frac{\rho+1}{\rho} \leqslant \frac{\rho+1}{\rho \cdot (\rho-1)} = 1.$$

Hence, by Theorem 34, greedy capture finds a $(1+\sqrt{2})$-proportionally fair clustering. □

Next, we show that for $d = L^2$, the $\rho$-Apollonius radius is slightly better, leading to a better 2-approximation guarantee for greedy capture.

**Theorem 36.** *For the metric space $(\mathbb{R}^t, L^2)$, where $t \in \mathbb{N}$, the $\rho$-Apollonius radius is $A_{\mathbb{R}^t,L^2}(\rho) \leqslant \frac{\rho}{\rho^2-1}$, and hence, greedy capture finds a 2-proportionally fair clustering.*

*Proof.* For the $L^2$ norm in a Euclidean space, it is well-known that given $x, y \in \mathbb{R}^t$ and $\rho > 1$, the set of points $\{p \in \mathbb{R}^t : \rho \cdot d(p,y) \leqslant d(p,x)\}$ is a ball of radius $d(x,y) \cdot \frac{\rho}{\rho^2-1}$. This is a simple algebraic exercise; its two-dimensional variant was known to Apollonius himself, after whom the result is named (the derivation is widely available online, e.g., see Cox and Partensky [2007]). This immediately implies that $A_{\mathbb{R}^t,L^2}(\rho) \leqslant \frac{\rho}{\rho^2-1}$.

Now, we have that

$$A_{\mathcal{X},d}(\rho) \cdot \frac{\rho+1}{\rho} \leqslant \frac{\rho}{\rho^2-1} \cdot \frac{\rho+1}{\rho} = \frac{1}{\rho-1}.$$

This quantity is at most 1 when $\rho$ is at least 2. Hence, by Theorem 34, greedy capture finds a 2-proportionally fair clustering for this metric. □

The obvious next question then is whether this refinement also provides an improved approximation bound for other distance metrics. Unfortunately, for two other prominent distance metrics, $L^1$ and $L^\infty$, the answer is no. We show this by providing a direct counterexample where greedy capture finds a clustering that is no better than $(1+\sqrt{2})$-proportionally fair.

**Theorem 37.** *For the metric space $(\mathbb{R}^t, d)$ where $t \geqslant 2$ and $d \in \{L^1, L^\infty\}$, and $\mathcal{M} = \mathbb{R}^t$, there exists an example in which the clustering produced by greedy capture is not $\rho$-proportionally fair for $\rho < 1+\sqrt{2}$.*

*Proof.* We provide an example with $t = 2$; for an example in higher dimension, we can simply set the remaining coordinates to 0. Further, for $\mathbb{R}^2$, we remark that $L^1$ and $L^\infty$ are equivalent[1] up to a rotation by $45\deg$. Hence, a counterexample for $L^1$, rotated by $45\deg$, is also a counterexample for $L^\infty$. Thus, let us use $d = L^1$ without loss of generality.

In our example, we have $|\mathcal{N}| = 28$ and $k = 7$. $\mathcal{N}$ consists of four isomorphic sets of 7 points each. The coordinates of one such set of points are given below.

[1]We say that two distance metrics $d_1$ and $d_2$ are equivalent if there exists a constant $\kappa > 0$ such that for all $x, y$, $d_1(x,y) = \kappa d_2(x,y)$. In this case, the equivalence is achieved after the space is rotated by $45\deg$.

| $p_1$ | $p_2$ | $p_3$ | $p_4$ | $p_5$ | $p_6$ | $p_7$ |
|---|---|---|---|---|---|---|
| $(0,1)$ | $(1,0)$ | $(0,-1)$ | $(-1,0)$ | $(0,\sqrt{2})$ | $(1+\epsilon,\sqrt{2})$ | $(-1-\epsilon,\sqrt{2})$ |

Each set is located sufficiently far from the other sets. Greedy capture opens the first four cluster centers at coordinates $(0,0)$ in each set. Then, it opens the remaining three cluster centers in three of the sets. Among points in the remaining fourth set, which derive their cost from the single center at $(0,0)$ in their set, points $\{p_2, p_5, p_6, p_7\}$ can deviate with a center at $(0,\sqrt{2})$, which would reduce the cost to each point by a factor at least $1+\sqrt{2}$ as $\epsilon \to 0$. Hence, the clustering produced by greedy capture is not $\rho$-proportionally fair for $\rho < 1+\sqrt{2}$. □

## 7.4 Universal Lower Bounds

In this section, we show lower bounds on approximation to proportional fairness that apply to all algorithms, as opposed to the lower bounds in the previous section that apply only to greedy capture. Chen et al. [2019] show that when $\mathcal{N}$, $\mathcal{M}$, and the metric are arbitrary, $\rho$-proportional fairness cannot be guaranteed for $\rho < 2$. They also consider the special case where $\mathcal{N} = \mathcal{M}$, and prove a slightly weaker lower bound of $1.5$. One question that they do not address is whether greedy capture provides better than $(1+\sqrt{2})$-approximation in this special case; As the example below shows, this is not the case.

**Example 9.** *Let $n = 45$ and $k = 9$. We consider an instance in which $\mathcal{N}$ consists of five isomorphic sets of $9$ points each. Each set is sufficiently far from all other sets. Let us describe one of the sets. Given five locations $x_1 < x_2 < x_3 < x_4 < x_5$ on a line, the set contains three points at $x_1$, three points at $x_5$, and one point at each of $x_2$, $x_3$, and $x_4$. The distances are as follows.*

| $d(x_1,x_2)=1$ | $d(x_2,x_3)=1$ | $d(x_3,x_4)=\sqrt{2}-1$ | $d(x_4,x_5)=1+\epsilon$ |
|---|---|---|---|

*Greedy capture first opens a center at $x_2$ in each set, and then four centers in at most four of the sets. Thus, in the remaining fifth set, all points derive their costs from a single center located at $x_2$. Then, the set of five points located at $x_3$, $x_4$, and $x_5$ can deviate (since $n/k = 5$) with $y = x_4$, which would reduce the cost to each by a factor of at least $1+\sqrt{2}$.*

In this section, we turn our attention to the case of our interest: $\mathcal{M} = \mathcal{X} = \mathbb{R}^t$ and $d \in \{L^1, L^2, L^{\inf}\}$. When $t = 1$, it is easy to notice that an exactly proportionally fair clustering always exists.[2] When $t \geqslant 2$, we provide a lower bound of $2/\sqrt{3}$ for $d = L^2$ and a lower bound of $1.4$ for $d \in \{L^1, L^\infty\}$.

**Theorem 38.** *For the metric space $(\mathbb{R}^t, L^2)$ where $t \geqslant 2$ and $\mathcal{M} = \mathbb{R}^t$, there is an example in which no clustering is $\rho$-proportionally fair for $\rho < 2/\sqrt{3} \approx 1.155$.*

*Proof.* Once again, we set $t = 2$ without loss of generality. Consider an instance in which $|\mathcal{N}| = 6$ and $k = 3$. Suppose $|\mathcal{N}|$ consists of two isomorphic sets of $3$ points each, where each set of $3$ points forms an equilateral triangle of length $1$ and the two sets are sufficiently far from each other. Then, by the pigeonhole principle, under any clustering $X$, at least one set of $3$ points, say $\{p_1, p_2, p_3\}$, must derive their costs from a single cluster center $x$.

Let $a$ denote the circumcenter of their triangle. Then, $d(a,p_1) = d(a,p_2) = d(a,p_3) = 1/\sqrt{3}$. Hence, $d(a,p_1) + d(a,p_2) + d(a,p_3) = \sqrt{3}$. Notice that in an equilateral triangle, the circumcenter

[2] For instance, opening a cluster at every $n/k$-th data point from left to right is proportionally fair.

is also the Fermat point, which minimizes the sum of distances from the three vertices. Hence, for the cluster center $x$, we have $d(x,p_1)+d(x,p_2)+d(x,p_3) \geqslant \sqrt{3}$. Without loss of generality, assume $d(x,p_1) \geqslant d(x,p_2) \geqslant d(x,p_3)$. Then, $d(x,p_1)+d(x,p_2) \geqslant 2/\sqrt{3}$.

Now, $p_1$ and $p_2$ can deviate, choose a location $y$ on the line joining $p_1$ and $p_2$ such that

$$d(y,p_1)/d(y,p_2) = d(x,p_1)/d(x,p_2).$$

Since $d(y,p_1)+d(y,p_2) = d(p_1,p_2) = 1$, this reduces the cost to each point by a factor of $2/\sqrt{3}$. Hence, the clustering is not $\rho$-proportionally fair for $\rho < 2/\sqrt{3}$. □

Note that the lower bound of $1.155$ is significantly lower than the upper bound of $2$ obtained by greedy capture for $L^2$ as shown in Theorem 36. Closing the gap is an interesting open question. Next, we show a lower bound for $L^1$ and $L^\infty$.

**Theorem 39.** *For the metric space $(\mathbb{R}^t, d)$, where $t \geqslant 2$ and $d \in \{L^1, L^\infty\}$, and $\mathcal{M} = \mathbb{R}^t$, there is an example in which no clustering is $\rho$-proportionally fair for $\rho < 1.4$.*

*Proof.* Once again, we set $t = 2$ without loss of generality, and using the earlier noted equivalence between $L^1$ and $L^\infty$ for $\mathbb{R}^2$ up to a rotation by $45\deg$, we only need to derive the lower bound for one of the two metrics. We do so for $L^\infty$ for simplicity of calculation.

Consider an instance in which $|\mathcal{N}| = 10$ and $k = 5$. Note that $n/k = 2$. Let $\mathcal{N}$ consist of two isomorphic sets of $5$ points each; the coordinates of one of the sets is given below. The other set, as usual, is assumed to be sufficiently far.

| $p_1$ | $p_2$ | $p_3$ | $p_4$ | $p_5$ |
|---|---|---|---|---|
| $(0,5)$ | $(1,0)$ | $(-1,0)$ | $(3,2)$ | $(-3,2)$ |

As before, by the pigeonhole principle, under any clustering $X$, at least one set of $5$ points (say the set described above, without loss of generality) must derive its costs from only two cluster centers.

Let those cluster centers by $c_1 = (x_1, y_1)$ and $c_2 = (x_2, y_2)$. Let $X = \{c_1, c_2\}$. Suppose for contradiction that the clustering is $\rho$-proportionally fair for some $\rho < 1.4$. We consider different cases regarding the centers' positions, and we prove that for each case, there exists a pair of points that can deviate to a new center $y$ and reduce the cost to each by a factor of at least $1.4$, establishing the contradiction.

First, we claim that $\min(x_1, x_2) < -0.6$. Indeed, if $\min(x_1, x_2) \geqslant -0.6$, then $d(p_3, X) \geqslant 0.4$ and $d(p_5, X) \geqslant 2.4$, while $d(p_3, p_5) = 2$. So, as in the proof of Theorem 38, they can choose a point $y$ on the line connecting them such that $d(y,p_3)/d(y,p_5) = d(p_3,X)/d(p_5,X)$, and reduce the cost to each by a factor of at least $1.4$. A symmetric argument with points $p_2$ and $p_4$ (instead of $p_3$ and $p_5$) shows that $\max(x_1, x_2) > 0.6$.

Next, we claim that $\max(y_1, y_2) > 1.4$. Indeed, if $\max(y_1, y_2) \leqslant 1.4$, then $d(p_1, X) \geqslant 3.6$ and $d(p_4, X) \geqslant 0.6$, while $d(p_1, p_4) = 3$. So, as before, they can choose a point on the line connecting them such that the cost to each reduces by a factor of at least $1.4$. A symmetric argument with points $p_2$ and $p_3$ (instead of $p_1$ and $p_4$) shows that $\min(y_1, y_2) < 1.4$.

Under these observations and as the instance is symmetric with respect to the $y$-axis, without loss of generality, we assume that $x_1 < -0.6$ and $y_1 > 1.4$, while $x_2 > 0.6$ and $y_2 < 1.4$. We distinguish between two cases.

- **Case 1:** $d(p_3, X) = d(p_3, c_1)$**.** Let $y_1 = 1.4 + y_1'$. Then, $d(p_3, X) \geqslant 1.4 + y_1'$. As $d(p_2, c_1) \geqslant 1.6$, we have that $d(p_2, X) = d(p_2, c_2) < 1.4 - y_1'$, otherwise $p_3$ and $p_2$ could choose a point on the line connecting them to reduce their costs by a factor of at least $1.4$. Hence, $x_2 < 2.4 - y_1$ and $y_2 < 1.4 - y_1$. Then, notice that $d(p_4, c_1) \geqslant 3.6$ and $d(p_4, c_2) \geqslant 0.6 + y_1'$, while $d(p_1, c_2) \geqslant 3.6$ and $d(p_1, c_1) \geqslant 3.6 - y_1'$. Now, it is easy to see that $p_1$ and $p_4$ can choose a center on the line connecting them and reduce their costs by a factor of at least $1.4$.

- **Case 2:** $d(p_3, X) = d(p_3, c_2)$**.** Let $x_2 = 0.6 + x_2'$. Then, $d(p_3, X) \geqslant 1.6 + x_2'$. As $d(p_5, c_2) \geqslant 3.6$, we must have $d(p_5, X) = d(p_5, c_1) < 1.2 - x_2'$. Hence, $x_1 < 1.8 - x_2'$ and $y_2 < 3.2 - x_2'$. Then, notice that $d(p_1, c_2) \geqslant 3.6$ and $d(p_1, c_1) \geqslant 5 - 3.2 + x_2'$, while $d(p_4, c_1) \geqslant 3.6$ and $d(p_4, c_2) \geqslant 2.4 - x_2'$. Now, we observe that $p_1$ and $p_4$ can choose a center on the line connecting them and reduce their costs by a factor of at least $1.4$.

This concludes the proof. □

## 7.5 Clustering in Graphs

In this section, we consider the special case where the metric space $(\mathcal{X}, d)$ is induced by an undirected graph $G = (V, E)$. Specifically, we let $\mathcal{X} = V$ be the set of nodes of the graph, and assume that $d(x, y)$ measures the length of the shortest path between nodes $x$ and $y$. As in the previous sections, we restrict our attention to the $\mathcal{M} = \mathcal{X}$ case, i.e., when every point of the metric space is a feasible cluster center.

Graphs are an important special case since clustering and facility location in graphs and networks are very well studied. However, while objectives such as truthfulness [Alon et al., 2010] and social welfare maximization (or social cost minimization) [Feldman and Wilf, 2013] have received significant attention, fairness has not.

We study proportional fairness for this setting. Our first result shows that when the graph $G = (V, E)$ is a tree, an exact proportionally fair clustering always exists, and can be computed by an efficient algorithm. Intuitively, the algorithm works as follows. We first root the tree at an arbitrary node $r$ to obtain a rooted tree $(G, r)$. We denote with $h$ the height of the rooted tree, and with $\mathrm{level}(x)$ the height of node $x$ relative to the root $r$ (with $\mathrm{level}(r) = 1$). Let $\mathrm{ST}(x)$ denote the subtree of node $x$ (i.e. the set of nodes $v$ such that with $\mathrm{level}(v) \geqslant \mathrm{level}(x)$ and the unique path from $v$ to $r$ contains $x$). The algorithm starts from the highest level (the leaves), opens a center every time it finds a node whose subtree contains at least $\lceil n/k \rceil$ nodes, and deletes this subtree from the graph. At the end, the cost to each node is still defined using the closest node at which a center is opened by the algorithm.

**Theorem 40.** *Let $G = (V, E)$ be an undirected tree, $(V, d)$ be the metric induced by $G$, $\mathcal{N} \subseteq V$, $\mathcal{M} = V$, and $k \in \mathbb{N}$. Then, Algorithm 6 yields a proportionally fair clustering.*

*Proof.* Let $X$ be the clustering returned by Algorithm 6. First, we notice that $X$ contains at most $k$ centers. This is because every time the algorithm opens a center in the for loop, it deletes at least $\lceil n/k \rceil$ nodes from the graph. Hence, the for loop can add at most $k$ centers. If it adds exactly $k$ centers, then the remaining graph must be empty, and the if condition is not executed. If it adds at most $k - 1$ centers, then even if the root node is added later, there would be at most $k$ centers.

Next, suppose for contradiction that $X$ is not proportionally fair. Hence, there exists a set $S \subseteq V$ with $|S| \geqslant \lceil n/k \rceil$ and $y \in V$ such that $d(i, y) < d(i, X)$ for all $i \in S$.

**Algorithm 6** Proportionally Fair Clustering for Trees

Root the tree $G$ at an arbitrary node $r$
$X \leftarrow \emptyset$
$G^h \leftarrow G$
**for** $\ell = h$ to $1$ **do**
  $G^{\ell-1} \leftarrow G^\ell$
  **for** every $x \in V$ with $\text{level}(x) = \ell$ and $|\operatorname{ST}(x)| \geqslant \lceil n/k \rceil$ **do**
    $X \leftarrow X \cup \{x\}$
    $G^{\ell-1} \leftarrow G^{\ell-1} \setminus \operatorname{ST}(x)$
  **end for**
**end for**
**if** $G^0 \neq \emptyset$ **then**
  $X \leftarrow X \cup \{r\}$
**end if**
**return** $X$

For each node $i \in V$, define $p(i)$ to be its closest ancestor in $X$ (i.e. $p(i) \in X$ and $i \in \operatorname{ST}(p(i))$, and $p(i)$ is the node of maximum level satisfying these two conditions). Note that because the algorithm always opens a center at the root node, $p(i)$ is well-defined for each node $i$.

Further, note that for each $i \in S$, $d(i, X) \leqslant d(i, p(i))$. And for nodes $j \notin \operatorname{ST}(p(i))$, $d(i, j) > d(i, p(i))$. Hence, the cost to $i$ can only reduce if the deviating center is in $\operatorname{ST}(p(i))$. We now consider two cases.

- **Case 1:** $\exists i, i' \in S : p(i) \neq p(i')$**.** First, suppose that $p(i)$ and $p(i')$ are siblings (i.e. $\operatorname{ST}(p(i)) \cap \operatorname{ST}(p(i')) = \emptyset$). As $i$ can only improve if the deviating center is in $\operatorname{ST}(p(i))$ and $i'$ can only improve if the deviating center is in $\operatorname{ST}(p(i'))$, we obtain that no deviating center $y$ can reduce the cost to $i$ and $i'$ simultaneously, which is a contradiction.

  Next, suppose that $p(i) \in \operatorname{ST}(p(i'))$. Then, we must have $y \in \operatorname{ST}(p(i))$, otherwise the cost to $i$ would not reduce. But then, $d(i', p(i)) \leqslant d(i', y)$. Hence, the cost to $i'$ does not reduce due to $y$, which is also a contradiction.

  The remaining case of $p(i') \in \operatorname{ST}(p(i))$ is symmetric to the last case.

- **Case 2:** $\forall i \in S, p(i) = p^*$**.** Let $O = X \cap \operatorname{ST}(p^*) \setminus \{p^*\}$ be the set of open centers in $\operatorname{ST}(p^*)$ except $p^*$ itself. Note that by definition of $p^*$, we have that if $i \in S$, then $i \notin \operatorname{ST}(o)$ for any $o \in O$.

  This implies that if $y \in \operatorname{ST}(o)$ for some $o \in O$, then for every point $i \in S$, we have $d(i, y) \geqslant d(i, o) \geqslant d(i, X)$, meaning that $y$ would not reduce the cost to any point in $S$. Hence, $y \in \operatorname{ST}(p^*) \setminus \cup_{o \in O} \operatorname{ST}(o)$.

  In other words, if center $p^*$ was opened in the iteration with index $\ell$ (let $\ell = 0$ if $p^*$ is the root node that was opened outside of the for loop), then $S \cup \{y\} \in G^\ell$ (any point from $S$ or $y$ could not have been deleted in any previous iteration). However, for $y$ to reduce the cost to each $i \in S$, we must have $S \subseteq \operatorname{ST}(y)$. However, then, $y$ is a node of higher level than $p^*$ that still contains at least $|S| \geqslant \lceil n/k \rceil$ points, so it must have been removed in a previous iteration. This is the desired contradiction.

This concludes the proof. □

We remark that Theorem 40 identifies the broadest class of interesting metrics to date for which an exact proportionally fair clustering is known to always exist. This raises an immediate question: what about graphs that are not trees? We can consider the universal lower bound for the $(\mathbb{R}^2, L^1)$ metric from Theorem 39. If we construct a very dense grid graph (in which the shortest path distance mimics the $L^1$ distance in the plane) in the relevant region of $\mathbb{R}^2$ from that example, we can derive the same lower bound of $1.4$ on the approximation ratio to proportional fairness for graphs. Whether better lower bounds exist is an open question. Similarly, the $1 + \sqrt{2}$ upper bound of greedy capture holds for all metrics, including those induced by a graph. However, whether its approximation ratio or running time can be improved for graphs, possibly with special properties such as planarity, remains to be seen.

It is worthwhile noting the special case of $\mathcal{N} = \mathcal{M} = V$, i.e., when every node of the graph is also a data point (besides being a feasible cluster center location). In this case, we do not know whether an exact proportionally fair clustering always exists, and leave this as an interesting open question. That said, we do note that if $G$ is connected and we want to place a large number of clusters $k \geqslant n/2$, then it can be shown that a proportionally fair clustering exists. This is because a dominating set[3] of any size $k \geqslant n/2$ is guaranteed to exist in a graph with $n$ nodes [Ore, 1962] and can be computed efficiently [Hedetniemi et al., 2015]. If nodes in such a set are chosen as the cluster centers, then every node in the graph already has cost at most $1$. So to deviate, all nodes in the blocking coalition must achieve cost $0$. However, since the blocking coalition must contain at least $\lceil n/k \rceil \geqslant 2$ nodes, this is impossible. Thus, the questions of existence and computation of a proportionally fair clustering in general graphs with $\mathcal{N} = \mathcal{M} = V$ become trivial when $k \geqslant n/2$, but remain open when $k < n/2$.

## 7.6 Computational Aspects

In this section, we consider computational aspects of two problems: the problem of checking whether a proportionally fair clustering exists, and the problem of implementing the greedy capture algorithm when $\mathcal{M} = \mathbb{R}^t$. We begin by considering the former problem.

**Theorem 41.** *Given $\mathcal{X} = \mathbb{R}^2$, finite $N \subset \mathcal{X}$, finite $M \subset \mathcal{X}$, $k \in \mathbb{N}$, and $d = L^2$, checking whether a proportionally fair clustering exists is NP-hard.*

*Proof.* We begin by proving the theorem for the case that $n/k = 2$, and later we generalize it for every case that $n/k$ is equal to an even integer.

We use a polynomial-time reduction from planar monotone rectilinear 3-SAT. Let $\phi$ be an instance of 3-SAT with $l$ boolean variables, $m$ clauses, and each clause is monotone, i.e. consists only of positive or negative literals. Let $I$ be an instance of a planar monotone rectilinear 3-SAT which consists of $\phi$. In a rectilinear embedding of planar 3-SAT, each variable is represented by a rectangle and all the rectangles are located on x-axis without intersections. We denote with $v_i$ the variable with the $i$th leftmost variable rectangle in $I$. Moreover, each clause is represented as a rectangle with three vertical lines or legs to its three literals. In the planar monotone rectilinear 3-SAT, all positive clauses (consisting only of positive literals), are drawn on the positive side of the variables and all negative clauses (consisting only of negative literals) are drawn on the negative side of the variables. We assume that there are $c_p$ positive

[3] A set of nodes is called a dominating set if every node in the graph is either in this set or adjacent to a node in this set.

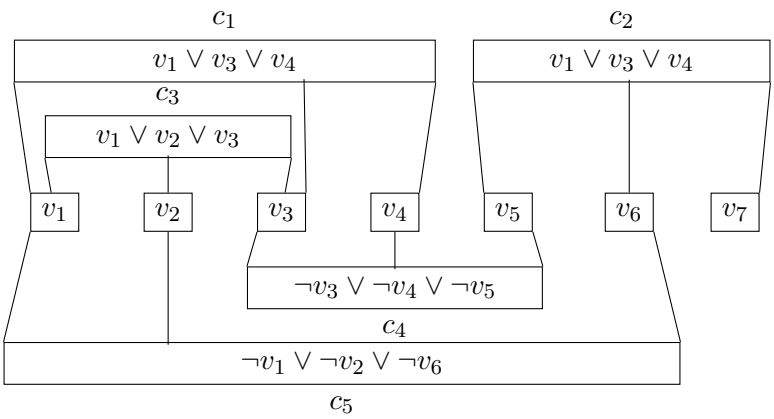


Figure 7.1: An example of labelling clauses and variables.

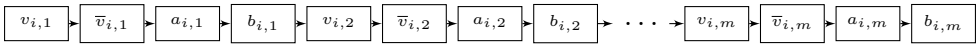


Figure 7.2: The variable gadget corresponding to variable $v_i$. Directed edges indicate the closest neighbour of each node.

and $c_n$ negative clauses (i.e. $c_p + c_n = m$). We label the clauses as following: start from the positive side and label with $c_j$ the clause with the $j$th largest $y$-value in $I$, and then continue with the negative side and label with $c_{j+c_p}$ the clause with the $j$th smallest $y$-value. We break ties from left to right. In figure 7.1, you can see a planar monotone rectilinear 3-SAT instance with the corresponding labelling of the variables and the clauses.

Let $\epsilon < \min\{\frac{2}{16ml-2}, \frac{1}{32lm^2+12m^2+8ml+4m}, \frac{\sqrt{2}-1}{4\sqrt{2}lm-1}, \frac{1}{8lm}, \frac{2.86}{32lm}\}$. Given $I$, we will construct an instance $I'$ of data points in 2-dimensions such that if $\phi$ is satisfiable then there is a proportionally fair clustering, otherwise there is no $(1+\epsilon)$-proportional solution. The high level idea is that we will construct an image similar with an instance of planar monotone rectilinear 3-SAT by placing points in the plane.

First, for each variable $v_i$, we construct a variable gadget which contains the nodes $v_{i,j}$, $\overline{v}_{i,j}$, $a_{i,j}$ and $b_{i,j}$, for $j \in [m]$. So, each variable gadget contains $4m$ points, and all these points are feasible cluster centers. The node $v_{i,j}$ (resp. $\overline{v}_{i,j}$) corresponds to the positive (resp. negative) literla of $v_i$, which may appear in clause $c_j$. All these nodes form a line as it is shown in figure 7.2. Regarding their distances, we set for every $j \in [m-1]$

- $d(v_{i,1}, \overline{v}_{i,1}) = 1 - (l-1)\epsilon - (4m-1)(i-1)\epsilon - \epsilon$
- $d(\overline{v}_{i,j}, a_{i,j}) = d(v_{i,j}, \overline{v}_{i,j}) - \epsilon$
- $d(a_{i,j}, b_{i,j}) = d(\overline{v}_{i,j}, a_{i,j}) - \epsilon$
- $d(v_{i,j+1}, \overline{v}_{i,j+1}) = d(b_{i,j}, v_{i,j+1}) - \epsilon$.

Intuitively, each node in a gadget has as closest neighbour the node on its right, and the distance of every two points in $x$-axis is at least equal to $1 - (l-1)\epsilon - (4m-1)l\epsilon = 1 - 4ml\epsilon + \epsilon > 1 - 4ml\epsilon$.

All the variable gadgets are located in x-axis in such a way that the gadget of variable $v_i$ is on the left of respective gadget of variable $v_{i+1}$ and we set $d(b_{i,m}, v_{i+1,1}) = 1 - i\epsilon > 1 - l\epsilon$, $\forall i \in [l-1]$. This

means that every point in a variable gadget prefers a point in its gadget than any point that belongs in a different one.

Then, we construct two clause gadgets, the basic and the auxiliary for each clause. We start by making a gadget which captures the rectangle and the legs of each clause in $I$ by placing some points. Let $y_j$ be the $y$-coordinate where we place all the points that consist the virtual rectangle of $c_j$. We set the values of $y_j$ as follows: for each positive clause $c_j$ we set $y_j = 8m + 1 - 8(j-1)$ and for each negative clause $c_j$ we set $y_j = -(8m + 1 - 8(j-1))$. Hence, the rectangles of two different clauses have distance at least equal to 8.

Let $c_j$ be a positive clause which contains the positive literals of variables $v_q$, $v_s$ and $v_t$ with $q < s < t$ (respectively, if the clause is negative). Our goal is to design a virtual rectangle in the interval $[(v_{q,j}, y_j), (v_{t,j}, y_j)]$ and add the three virtual legs that are vertical to $(v_{q,j}, 0)$, $(v_{s,j}, 0)$ and $(v_{t,j}, 0)$, respectively.

We start from the virtual rectangle. Notice that in the interval $[(\overline{v}_{s,j}, 0), (v_{t,j}, 0)]$, there are $2m(t-s)$ points. We place $2m(t-s)$ points in the interval $[(\overline{v}_{s,j}, y_j), (v_{t,j}, y_j)]$. More specifically, we add one point in position $(a_{s,j}, y_j)$ and then every two points in $x$-axis we add one point at the respective $x$-coordinate with $y$-coordinate equal to $y_i$. Denote these points as $r^1_{j,k}$, $k \in [1, 2m(t-s)]$, where the $x$-coordinate of $r^1_{j,k}$ is less than the $x$-coordinate of $r^1_{j,k+1}$. At the end, given the distances of the points that are located in the interval $[(a_{s,j}, 0), (v_{t,j}, 0)]$, we shift the points $r^1_{j,k}$, with $k \in [2, 2m(t-s)-1]$ in such a way that every node has as closest neighbour the point on its right. More precisely, let $\vec{\sigma}$ be a vector that contains all the distances of the points in $[(a_{s,j}, 0), (v_{t,j}, 0)]$ in a decreasing order, then we set $d(r^1_{j,1}, r^1_{j,2}) = \sigma(1) + \sigma(2)$, $d(r^1_{j,2}, r^1_{j,3}) = \sigma(3) + \sigma(4)$ and so on.

Next, we place in a similar way $2m(s-q)$ points in the interval $[(v_{q,j}, y_j), (a_{s,j-1}, y_j)]$ (or if $j = 1$, in the interval $[(v_{qj}, y_j), (a_{s-1,m}, y_j)]$ ). We denote these points as $r^2_{j,k}$, $k \in [1, 2m(s-q)]$, where the $x$-coordinate of point $r^2_{j,k}$ is larger than the $x$-coordinate of point $r^2_{j,k+1}$ and $d(r^2_{j,k}, r^2_{j,k+1}) > d(r^2_{j,k+1}, r^2_{j,k+2})$.

Notice that $d(r^1_{j,1}, r^2_{j,1}) > 4 - 16ml\epsilon$, while $d(r^1_{j,1}, r^1_{j,2}) < 2 - 2\epsilon$ and $d(r^2_{j,1}, r^2_{j,2}) < 2 - 2\epsilon$. Thus, by the definition of $\epsilon$ we have that the closest neighbour of $r^1_{j,1}$ (resp. $r^2_{j,1}$ ) is $r^1_{j,2}$ (resp. $r^2_{j,2}$ ). Lastly, as mentioned two points that belong in different rectangles have distance at least 8 and so, the closest neighbour of every $r^1_{j,k}$ (resp. $r^2_{j,k}$ ) is $r^1_{j,k+1}$ (resp. $r^2_{j,k+1}$).

It remains to construct the legs of each clause. First, we place $4m - 4(j-1)$ points, denoted by $l^1_{j,k}$, $k \in [1, 4m - 4(j-1)]$, with $x$-coordinate equal to $v_{q,j}$ and $y$-coordinate less than $8m + 1 - 8(j-1)$. Specifically, we locate $l^1_{j,1}$ in a position such that its distance from the leftmost point in the rectangle of $c_j$ is equal to $2 - 8ml\epsilon$ and then we place all the points in such a way that: $d(l^1_{j,k}, l^1_{j,k+1}) = d(l^1_{j,k-1}, l^1_{j,k}) - \epsilon$. Hence, every two points in this leg have distance at least equal to $2 - 8ml\epsilon - 4m\epsilon$. Notice that the closest neighbour of $r^2_{j,2m(s-q)}$ is $l^1_{j,1}$, while the closest neighbour of every $l^1_{j,k}$ with $k \in [1, 4m - 4(j-1) - 1]$ is $l^1_{j,k+1}$.

Now, we will show that the closest neighbour of $l^1_{j,4m-4(j-1)}$ is $v_{q,j}$. Instead, as every two points in this leg have distance at least equal to $2 - 8ml\epsilon - 4m\epsilon$ if we add all the distances of the points in the leg, we conclude in a line with length at least equal to

$$\sum_{k=1}^{4m-4(j-1)} 2 - 8ml\epsilon - 4m\epsilon = 8m - 8(j-1) - (4m - 4(j-1))(8ml\epsilon + 4m\epsilon) \geqslant$$
$$4m - 4(j-1) - (32lm^2 + 16m^2)\epsilon.$$

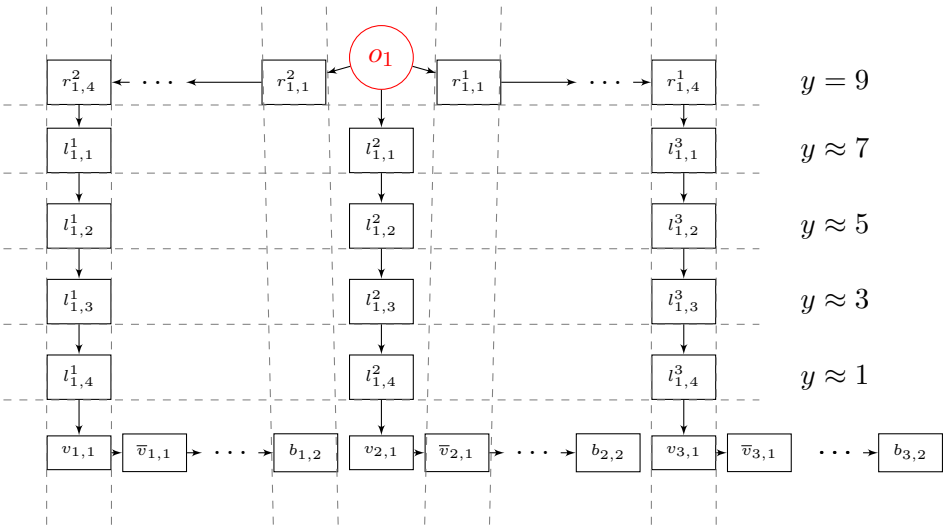


Figure 7.3: The rectangle and the legs of clause $c_1 = v_1 \vee v_2 \vee v_3$ in an instance with $m = 2$ and $l = 3$. A directed edge from one node points to the closest neighbor of the node.

Hence,

$$1 + (32lm^2 + 16m^2)\epsilon \geqslant d(l^1_{j,4m-4(j-1)}, v_{q,j}) \geqslant 1$$

and from the definition of $\epsilon$, we obtain that $2 - 8ml\epsilon - 4m\epsilon > 1 + (32lm^2 + 16m^2)\epsilon$. However, $v_{q,j}$ still prefers $\overline{v}_{q,j}$ as $d(v_{q,j}, \overline{v}_{q,j}) < 1 - l\epsilon$.

In a similar way, we construct the two remaining legs. More specifically, we add points at exactly the same $y$-coordinates of the points in the first leg, but with $x$-coordinates equal to $v_{s,j}$ and $v_{t,j}$. Denote the points of the middle and the right leg as $l^2_{j,k}$ and $l^3_{j,k}$, $k \in [1, 4m - 4(j-1)]$, respectively.

Now, we claim that $d(r^1_{j,1}, r^1_{j,2}) < d(r^1_{j,1}, l^2_{j,1})$. Instead, we already know that $d(r^1_{j,1}, (u_{s,j}, y_j)) > 2 - 8lm\epsilon$ and $d(l^2_{j,1}, (u_{s,j}, y_j)) = 2 - 8lm\epsilon$, and so from Pythagorean theorem we have that $d(r^1_{j,1}, l^2_{j,1}) > \sqrt{2} \cdot (2 - 8lm\epsilon)$. Hence, from the definition of $\epsilon$, we have that $d(r^1_{j,1}, r^1_{j,2}) < 2 - 2\epsilon < \sqrt{2} \cdot (2 - 8lm\epsilon)$. With similar arguments, we conclude that $d(r^2_{j,1}, r^2_{j,2}) < d(r^2_{j,1}, l^2_{j,1})$.

At this point, notice that each clause gadget has an even number of points, denoted by $2n_j$, and these points are at most $12m + 2m(l-1)$. All the points that have been added are feasible cluster centres, too. Now, for each clause we add one more point, denoted by $o_j$ in a location such that it is the circumcenter of the triangle with nodes $r^1_{j,1}$, $r^2_{j,1}$ and $l^2_{j,1}$, i.e. $d(o_j, r^1_{j,1}) = d(o_j, r^2_{j,1}) = d(o_j, l^2_{j,1})$.

The next lemma shows that $o_j$ has as closest neighbours the points $r^1_{j_1}$, $r^2_{j,1}$ and $l^2_{j,1}$.

**Lemma 20.** *The closest neighbours of $o_j$ are $r^1_{j,1}$, $r^2_{j,1}$ and $l^2_{j,1}$*

*Proof.* Consider the right triangle that is formed from the vertices $r^1_{j,1}$, $(v_{s,j}, y_j)$ and $l^2_{j,1}$. We know that $d(r^2_{j,1}, (v_{s,j}, y_j)) > d(l^2_{j,1}, (v_{s,j}, y_j))$ and from the definition of $\epsilon$, it holds that $d(l^2_{j,1}, (v_{s,j}, y_j)) = 2 - 8lm\epsilon > 1$, while $d(r^2_{j,1}, (v_{s,j}, y_j)) < 2 - 2\epsilon < 2$. Hence, the angle $\angle r^1_{j,1} l^2_{j,1} (v_{s,j}, y_j)$ is less than $64°$ and more than $45°$. Now, consider the right triangle that is formed from the vertices $r^2_{j,1}$, $(v_{s,j}, y_j)$ and $l^2_{j,1}$. With similar arguments, we obtain that the angle $\angle r^2_{j,1} l^2_{j,1} (v_{s,j}, y_j)$ is less than $64°$ and more than

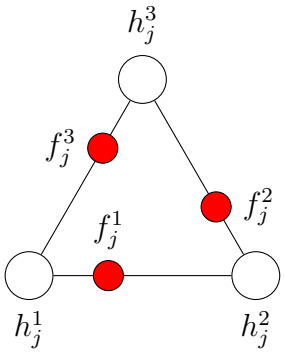


Figure 7.4: The auxiliary clause gadget corresponding to a clause $c_j$.

$45°$. Thus, in the triangle that is formed by the vertices $r_{j,1}^1$, $r_{j,1}^2$ and $l_{j,1}^2$, the angle $\angle r_{j,1}^1 l_{j,1}^2 r_{j,1}^1$ is less than $128°$, and more than $90°$. It is known that the diameter of the circumcircle, can be computed as the length of any side of the triangle divided by the sine of the opposite angle. Thus, the diameter of the circumcircle, with circumcenter $o_j$ is at most equal to $d(r_{j,1}^2, r_{j,1}^2)/sin(128°) < 4/0.78$, and hence $d(o_j, r_{j,1}^1) < 2.57$.

Now, the distance of $o_j$ to the points of a different rectangle is at least $8 - 2.57 > 2.57$. Moreover, the distance of $o_j$ to the points of a leg of a different clause is at least equal to $8 - 32l\epsilon - 2.57 > 2.57$, where the last inequality follows from the definition of $\epsilon$. □

Lastly, for each clause we construct an extra auxiliary clause gadget as shown in figure 7.4. More specifically, we add three points $h_j^1$,$h_j^2$ and $h_j^3$ which form an equilateral triangle with sides equal to 1, and they are not feasible cluster centers. The red spots, $f_j^1$, $f_j^2$ and $f_j^3$, are feasible cluster centers. The $f_j^1$ point is located on the edge of $h_j^1$ and $h_j^2$ in a position such that $d(h_j^1, f_j^1) = 0.25$ and $d(h_j^2, f_j^1) = 0.75$. The $f_j^2$ point is located on the edge of $h_j^2$ and $h_j^3$ in a position such that $d(h_j^2, f_j^2) = 0.25$ and $d(h_j^3, fs_j^1) = 0.75$. Lastly, the point $f_j^3$ is located on the edge of $h_j^1$ and $h_j^3$ in a position such that $d(h_j^3, f_j^3) = 0.25$ and $d(h_j^1, f_j^3) = 0.75$.

Hence, we conclude in an instance in which each clause gadget, which contains the basic and the auxiliary gadget, has an even number of points equal to $2n_j + 4$, and each variable gadget has $4m$ of points. So, we choose $k = 2ml + \sum_{j=1}^{m} n_j + 2m$, and hence $n/k = 2$. Now, we are ready to prove that if $\phi$ is satisfiable, then there is a proportional solution in $I'$ and vice versa.

First, we show that if $\phi$ is satisfiable, then we can use the assignment of the variables to find a proportional solution. Notice that in each variable gadget we need at least $2m$ points in order to construct a proportional clustering, as every two consecutive points could deviate if none of them is an open center. There are only two ways to place exactly $2m$ centers: the first choice is to open as centers the points $v_{i,j}$ and $a_{i,j}$ and the second choice is to open the points $\overline{v}_{i,j}$ and $b_{i,j}$. The first way corresponds to the assignment of true to $x_i$, while the second corresponds to the assignment of false to $x_i$.

Next, we observe that the auxiliary gadget of each clause always needs two cluster centers to achieve a proportional clustering. So, $n_j$ centers remain for every basic clause gadget. Moreover, we see that each basic gadget consumes at least $n_j$ centers. To see this notice that starting from point $r_{j,1}^2$ till the last point of the left leg, we need to add one center every two points. Similarly, starting from point $r_{j,1}^1$ till the last point of the right leg. In addition, in the middle leg we consume one center every two points. As all these points are $2n_j$, we need to place $n_j$ centers. However, at least one of $r_{j,1}^1$, $r_{j,1}^2$ and $l_{j,1}^2$ should

be an open center, otherwise $o_j$ has an incentive to choose one of them as center. This is possible only if at least one of the corresponding variable nodes is an open center. The reason is that the last point of at least one leg is not an open center and this does not have an incentive to deviate with its closest node (which corresponds to the literal of the clause) if and only if this node is an open center, and so the clause is satisfied.

From the above discussion immediately follows the opposite direction.

Regarding the generalization of the statement, if we replace each point with $r$ points in the same position, then $n/k = 2r$, and this instance is equivalent with the case that in each position there is only one point and $n/k = 2$. □

Next, we consider implementing the greedy capture algorithm when $\mathcal{M} = \mathbb{R}^t$. Note that because the naive description of greedy capture requires simultaneously growing a ball from each point in $\mathcal{M}$, it is easy to implement when $\mathcal{M}$ is finite (as shown by Chen et al. [2019]) but difficult when $\mathcal{M}$ is infinite.

To avoid this issue, one may consider restricting the set of feasible cluster center locations to a finite set $\mathcal{M}'$, and hope that running greedy capture (or any algorithm for that matter) on this finite set still yields a clustering that is approximately fair with respect to the original infinite $\mathcal{M}$. To that end, we show that restricting to $\mathcal{M}' = \mathcal{N}$ (i.e. choosing cluster centers from the set of data points) can only worsen the approximation factor of an algorithm by a factor of at most 2.

**Theorem 42.** *For any metric space $(\mathcal{X}, d)$, $k \in \mathbb{N}$, $\mathcal{N}, \mathcal{M} \subseteq \mathcal{X}$, and $\rho \geqslant 1$, any $\rho$-proportionally fair $k$-clustering with respect to $(\mathcal{N}, \mathcal{M}' = \mathcal{N})$ is $2\rho$-proportionally fair with respect to $(\mathcal{N}, \mathcal{M})$.*

*Proof.* Let $(\mathcal{X}, d)$ be any metric space, and $\mathcal{N}, \mathcal{M} \subseteq \mathcal{X}$. For some $\rho \geqslant 1$, let $X$ be a $k$-clustering that is $\rho$-proportionally fair with respect to $(\mathcal{N}, \mathcal{M}' = \mathcal{N})$.

Consider any set of points $S \subseteq \mathcal{N}$ with $|S| \geqslant n/k$, and any feasible cluster center $y \in \mathcal{M}$. Let $i^* \in \arg\min_{i \in S} d(i, y)$ be the closest point in $S$ from $y$. Because $X$ is $\rho$-proportionally fair with respect to $(\mathcal{N}, \mathcal{M}' = \mathcal{N})$, and $i^* \in \mathcal{N}$, there must exist $i \in S$ such that

$$d(i, X) \leqslant \rho \cdot d(i, i^*) \leqslant \rho \cdot (d(i, y) + d(y, i^*)) \leqslant 2\rho \cdot d(i, y),$$

where the second transition follows from the triangle inequality, and the third transition holds the definition of $i^*$ and the fact that $i, i^* \in S$ imply $d(i^*, y) \leqslant d(i, y)$.

Thus, $X$ must be $2\rho$-proportionally fair with respect to $(\mathcal{N}, \mathcal{M})$. □

For the metric $(\mathbb{R}^t, L^2)$, where $t \in \mathbb{N}$, recall that greedy capture produces a 2-proportionally fair clustering when $\mathcal{M} = \mathbb{R}^t$. Hence, Theorem 42 implies that running it on $\mathcal{M}' = \mathcal{N}$, for finite $\mathcal{N}$, would yield a 4-proportionally fair clustering in polynomial time. Moreover, Aziz et al. [2023b] show that there exists a different polynomial time algorithm that returns a 3-proportionally fair clustering for any $d$. However, for this special case where $d = L^2$, we can in fact efficiently achieve $2 + \epsilon$ approximation for any constant $\epsilon > 0$, as we discuss below.

Let us consider the first cluster center added by greedy capture when $\mathcal{M} = \mathbb{R}^t$. It is the center of a smallest ball that contains at least $n/k$ of the $n$ given points. The problem of finding the smallest ball containing at least $p$ of $n$ given points is a well-studied problem in computational geometry. This is known to be NP-hard [Shenmaier, 2013], which is what makes implementing greedy capture hard when $\mathcal{M} = \mathbb{R}^t$.

However, when $t$ is constant, the problem is known to be solvable efficiently [Shenmaier, 2013]; we show that in this case, greedy capture can also be implemented efficiently, and we can achieve the 2-approximation from Theorem 36.

Further, even when $t$ is large, the problem admits a PTAS [Shenmaier, 2013]. We show that this PTAS can be used to approximately implement greedy capture, and the approximation ratio only slightly worsens.

**Theorem 43.** *Let $t \in \mathbb{N}$, finite $\mathcal{N} \subset \mathbb{R}^t$, and $k \in \mathbb{N}$ be given as input. Suppose $\mathcal{M} = \mathbb{R}^t$ and $d = L^2$. Then, the following hold.*

1. *The clustering returned by greedy capture algorithm cannot be computed in polynomial time unless $P = NP$.*
2. *If $t$ is constant, then it can be computed in polynomial time.*
3. *Even if $t$ is not constant, for any constant $\epsilon > 0$, there exists a polynomial-time algorithm which finds a $(2+\epsilon)$-proportionally fair clustering.*

*Proof.* The $p$-minimum enclosing ball problem consists of $t \in \mathbb{N}$, a set $\mathcal{N}$ of $n$ points in $\mathbb{R}^t$, and $p \leqslant n$ as input, and the goal is to compute the radius of the smallest ball that contains at least $p$ points. We note that this problem is NP-hard in general, but efficiently solvable when $t$ is constant [Shenmaier, 2013].

For the first part, we reduce this problem to the problem of implementing greedy capture. In particular, we choose $k$ such that $p = \lceil n/k \rceil$, and run the greedy capture algorithm. Let $X$ be the set of cluster centers it returns. Note that the center of the smallest ball that contains at least $p$ points must be in $X$, as it must be the first location where greedy capture opens a center. Thus, from each point $x \in X$, we compute the distances to all points in $\mathcal{N}$ and compute the $p$-th smallest distance. This is precisely the smallest radius of a ball centered at $x$ that contains at least $p$ points. Minimizing this over $x \in X$ yields the smallest radius of any ball that contains at least $p$ points, and thus solves the $p$-minimum enclosing ball problem.

For the second part, we note that implementing greedy capture simply requires iteratively solving the $\lceil n/k \rceil$-minimum enclosing ball problem. In particular, once we solve the problem in a given iteration, and if the radius returned is $r$, we can enlarge the balls at previously opened centers to have radius $r$ and check if they would cover any more points before we can open the smallest new ball containing at least $\lceil n/k \rceil$ points. If they do, then we remove such points and then re-solve the $\lceil n/k \rceil$-minimum enclosing ball problem. Note that each such iteration removes at least one point, and hence, we do not need to solve $\lceil n/k \rceil$-minimum enclosing ball problem more than $n$ times. Hence, greedy capture can be implemented efficiently when $t$ is constant.

For the third part, we note that there exists a PTAS for the $\lceil n/k \rceil$-minimum enclosing ball problem [Shenmaier, 2013]. That is, given any constant $\epsilon > 0$, there exists a polynomial-time algorithm that is guaranteed to find a ball of radius at most $R \cdot (1+\epsilon)$ containing at least $\lceil n/k \rceil$ points, if $R$ is the smallest radius of any such ball. We note that running this subroutine instead of an exact algorithm for $\lceil n/k \rceil$-minimum enclosing ball problem approximately preserves the approximation guarantee derived in Theorem 36. In particular, in the proof of Theorem 34, we can guarantee that $d(i,x) \leqslant A_{\mathcal{X},d}(\rho) \cdot d(x,y) \cdot (1+\epsilon)$. This establishes that this variant of greedy capture provides $\rho$-approximation if $\rho$ satisfies $A_{\mathcal{X},d}(\rho) \cdot \frac{\rho+1}{\rho} \cdot (1+\epsilon) \leqslant 1$. Substituting the bound on $A_{\mathbb{R}^t,L^2}(\rho)$ from Theorem 36, we get that the smallest $\rho$ satisfying this equation is $2+\epsilon$, yielding the desired result. □

## 7.7 Learning Fair Clustering

A key concern in machine learning is generalization. In our context, the question is whether a clustering that is (approximately) proportionally fair with respect to random samples taken from an underlying population would remain (approximately) proportionally fair with respect to the whole population. A positive answer to this question could be useful in two ways.

First, sometimes we may have access only to samples from an underlying population. In this case, we can rest assured that by computing a clustering that is fair with respect to the samples, it is also fair with respect to the population. Second, even if the entire population is known, it may be very large. As we noticed in Section 7.6, finding a proportionally fair clustering or even running the greedy capture algorithm is NP-hard; thus, these tasks may be infeasible for a large population. However, it may be possible to do so on a smaller sample taken from the population, which is where the generalization guarantee can be useful.

Chen et al. [2019] show that generalization indeed holds for proportional fairness. Specifically, they define the following relaxation of $\rho$-proportional fairness.

**Definition 15.** *We say that a $k$-clustering $X$ is $\rho$-proportionally fair to $(1+\epsilon)$-deviations with respect to $\mathcal{N}$ if for all $S \subseteq \mathcal{N}$ with $|S| \geqslant |\mathcal{N}| \cdot (1+\epsilon)/k$ and all $y \in \mathcal{M}$, there exists at least one $i \in S$ such that $\rho \cdot d(i,y) \geqslant d(i,X)$.*

Chen et al. show that if $N \subseteq \mathcal{N}$ is a uniformly random sample of size $|N| = \Omega\left(\frac{k^3}{\epsilon} \ln \frac{|\mathcal{M}|}{\delta}\right)$, and if $X$ is $\rho$-proportionally fair with respect to $N$, then $X$ is $\rho$-proportionally fair to $(1+\epsilon)$-deviations with respect to $\mathcal{N}$ with probability at least $1-\delta$.

Unfortunately, this bound depends on $|\mathcal{M}|$, and breaks down when $|\mathcal{M}|$ is infinite, which is the focus of our work. We establish a stronger guarantee that does not depend on $|\mathcal{M}|$ by utilizing the framework of VC dimension [Vapnik and Chervonenkis, 1971] for binary classifiers. First, we show that there is a natural family of binary classifiers associated with a given clustering.

**Definition 16.** *Given a set of points $\mathcal{N}$, a $k$-clustering $X \in [\mathcal{M}]^k$, and $y \in \mathcal{M}$, define the binary classifier $h_{X,y} : \mathcal{N} \to \{0,1\}$ such that $h_{X,y}(i) = 1$ if and only if $\rho \cdot d(i,y) < d(i,X)$. Define the "error" of this classifier on a set of points $S \subseteq \mathcal{N}$ as $\mathrm{err}_S(h_{X,y}) = (1/|S|) \cdot \sum_{i \in S} h_{X,y}(i)$.*

Intuitively, $h_{X,y}(i) = 1$ if and only if $i$ can be part of a coalition that complains about the unfairness of $X$ by demonstrating $y$ as a location that provides them $\rho$-improvement. The use of the term "error" may be confusing. Unlike in traditional classification context, where there is a true classifier and the error is measured in terms of the fraction of points on which a given classifier differs from the true classifier, in our case the "error" is simply the fraction of points that can deviate. One can equivalently think of the "true classifier" as the one that outputs 0 on every point.

Note that $X$ is $\rho$-proportionally fair to $(1+\epsilon)$-deviations with respect to $\mathcal{N}$ if and only if $\mathrm{err}_{\mathcal{N}}(X,y) \leqslant \frac{1+\epsilon}{k}$ for all $y \in \mathcal{M}$. Our goal is to show that given a sufficiently large random sample $N \subseteq \mathcal{N}$, if we have a clustering $X$ that is $\rho$-proportionally fair with respect to $N$, then it is $\rho$-proportionally fair to $(1+\epsilon)$-deviations with respect to $\mathcal{N}$ with high probability. However, note that we have no control over what $X$ or $y$ are. This is where the stronger "uniform convergence" guarantee — this establishes that $|\mathrm{err}_N(h_{X,y}) - \mathrm{err}_{\mathcal{N}}(h_{X,y})|$ is bounded for all $X, y$ — becomes useful. Let us begin by introducing the VC dimension and a well-known uniform convergence guarantee that depends on the VC dimension.

**Definition 17** (VC Dimension)**.** *Let $\mathcal{N}$ be a set of points. Let $\mathcal{H}$ be a family of binary classifiers over $\mathcal{N}$. We say that $\mathcal{H}$* shatters *$S \subseteq \mathcal{N}$ if for every labeling $\ell : S \to \{0,1\}$, there exists a classifier $h \in \mathcal{H}$ such that $h(i) = \ell(i)$ for all $i \in S$. The VC dimension of $\mathcal{H}$, denoted $\dim^{\mathrm{VC}}(\mathcal{H})$, is the size of the largest $S$ that can be shattered by $\mathcal{H}$.*

**Proposition 7** ([Shalev-Shwartz and Ben-David, 2014])**.** *Let $\mathcal{H}$ be a family of binary classifiers over a set of points $\mathcal{N}$. If $N \subseteq \mathcal{N}$ is a uniformly random sample with $|N| \geqslant \Omega\left(\frac{\dim^{\mathrm{VC}}(\mathcal{H})+\ln(1/\delta)}{\epsilon^2}\right)$, then with probability at least $1-\delta$, $|err_N(h) - err_{\mathcal{N}}(h)| \leqslant \epsilon$ for all $h \in \mathcal{H}$.*

We show that the family of classifiers $\left\{h_{X,y} | X \in [\mathcal{M}]^k, y \in \mathcal{M}\right\}$ has finite VC dimension when $\mathcal{M} = \mathbb{R}^t$ with finite $t$. This, along with Proposition 7, gives us the desired result.

**Theorem 44.** *Fix $\epsilon, \delta > 0$, $\rho \geqslant 1$, $k, t \in \mathbb{N}$, and metric $(\mathbb{R}^t, d)$ where $d = L^2$. Let $\mathcal{N}$ be a set of points and $\mathcal{M} = \mathbb{R}^t$ be the set of feasible cluster centers. Let $N \subseteq \mathcal{N}$ be sampled uniformly at random with $|N| \geqslant \Omega\left(\frac{k^2 \cdot (tk \ln k + \ln(1/\delta))}{\epsilon^2}\right)$. Then, with probability at least $1-\delta$, every $k$-clustering $X \in [\mathcal{M}]^k$ that is $\rho$-proportionally fair with respect to $N$ is $\rho$-proportionally fair to $(1+\epsilon)$-deviations with respect to $\mathcal{N}$.*

*Proof.* Given a pair of points $x$ and $y$, note that the set of points $i$ such that $\rho \cdot d(i,y) \geqslant d(i,x)$ is a half-space in $\mathbb{R}^t$ when $\rho = 1$ and a ball in $\mathbb{R}^t$ when $\rho > 1$. Hence, given $X \in [\mathcal{M}]^k$, the set of points $i$ satisfying $\rho \cdot d(i,y) \geqslant d(i,X)$ is the union of $k$ half-spaces or balls in $\mathbb{R}^t$. It is known that the VC dimension of unions of $k$ half-spaces or balls in $\mathbb{R}^t$ is $O(tk \ln k)$ [Blumer et al., 1989].

Substituting this bound in Proposition 7, we get that $|N| \geqslant \Omega\left(\frac{k^2 \cdot (tk \ln k + \ln(1/\delta))}{\epsilon^2}\right)$ is sufficient to ensure that with probability at least $1-\delta$, $\mathrm{err}_{\mathcal{N}}(h_{X,y}) \leqslant \mathrm{err}_N(h_{X,y}) + \epsilon/k$ for all $X, y$. In particular, when $X$ is $\rho$-proportionally fair with respect to $N$, this ensures that with probability at least $1-\delta$, $X$ is $\rho$-proportionally fair to $(1+\epsilon)$-deviations with respect to $\mathcal{N}$. □

While we do not formally consider the case of infinite set of points ($|\mathcal{N}| = \infty$) in this work, one can define $\mathcal{N}$ as a distribution over infinitely many points, and ask whether a probability mass of at least $1/k$ has a beneficial deviation. Note that Theorem 44 applies to this case as well because it does not depend on $|\mathcal{N}|$. On the other hand, note that Theorem 44 applies only for $L^2$ distance metric; deriving generalization bounds for other distance metrics remains an interesting challenge for future work.

## 7.8 Discussion

In this work, we advanced the study of proportionally fair clustering in a metric space, and focused on the case where the set of possible cluster center $\mathcal{M}$ is the entire (usually infinite) metric space. Our work leaves a number of open questions.

The most immediate question is to bridge the gap between our lower and upper bounds on the approximation ratio to proportional fairness from Sections 7.3 and 7.4. In particular, we believe that for $L^2$, the lower bound of $2/\sqrt{3}$ from Theorem 38 may be achievable. This specific number is reminiscent of Jung's theorem, which states that for $L^2$ distance in $\mathbb{R}^2$ (which is where the lower bound stems from), any set of points with diameter at most 1 is contained in a ball of radius at most $1/\sqrt{3}$. This could be useful in closing the gap for $L^2$. Section 7.5 leaves open an important question which is whether a proportionally fair clustering exists for all graphs when $\mathcal{N} = \mathcal{M}$ is the set of all nodes of the graph. Our

hardness results from Section 7.6 and learnability results from Section 7.7 only apply to the $L^2$ distance because they use results and techniques from the literature that are only available for $L^2$. Deriving similar results for other distance metrics would be very interesting. While Chen et al. [2019] show that achieving bounded approximation to proportional fairness is incompatible with achieving bounded approximation to three classic objective functions ($k$-center, $k$-means, and $k$-median), it would be interesting to study the tradeoff between fairness and these objectives in special cases (such as graphs).

In this chapter, our focus was on the case where the set of possible cluster centers $\mathcal{M}$ is the entire metric space. In the case where the metric is induced by a graph $G = (V, E)$ (Section 7.5), it is also interesting to consider the case where $\mathcal{M} \subset V$. Does a proportionally fair clustering always exist for trees when $\mathcal{M} \subset V$, or for general graphs when $\mathcal{N} = \mathcal{M} \subset V$? These questions remain open.

More broadly, we find proportional fairness to be a very elegant fairness solution concept for clustering. Adapting this idea of proportional fairness to other machine learning settings such as regression or classification can lead to many avenues for future work.

# Chapter 8

# Fair Algorithms for Multi-Agent Multi-Armed Bandits

## 8.1 Introduction

In the classical (stochastic) multi-armed bandit (MAB) problem, a principal has access to $K$ arms and pulling arm $j$ generates a stochastic reward for the principal from an unknown distribution $D_j$ with an unknown mean $\mu_j^*$. If the mean rewards were known a priori, the principal could just repeatedly pull the *best arm* given by $\arg\max_j \mu_j^*$. However, the principal has no apriori knowledge of the quality of the arms. Hence, she uses a learning algorithm which operates in rounds, pulls arm $j^t$ in round $t$, observes the stochastic reward generated, and uses that information to learn the best arm over time. The performance of such an algorithm is measured in terms of its cumulative regret up to a horizon $T$, defined as $\sum_{t=1}^{T}(\max_j \mu_j^* - \mu_{j^t}^*)$. Note that this is the difference between the total mean reward that would have been achieved if the best arm was pulled repeatedly and the total mean reward of the arms pulled by the learning algorithm up to round $T$.

This problem can model situations where the principal is deliberating a policy decision and the arms correspond to the different alternatives she can implement. However, in many real-life scenarios, making a policy decision affects not one, but several agents. For example, imagine a company making a decision that affects all its employees, or a conference deciding the structure of its review process, which affects various research communities. This can be modeled by a multi-agent variant of the multi-armed bandit (MA-MAB) problem, in which there are $N$ agents and pulling arm $j$ generates a (possibly different) stochastic reward for each agent $i$ from an unknown distribution $D_{i,j}$ with an unknown mean $\mu_{i,j}^*$.

Before pondering about learning the "best arm" over time, we must ask what the best arm even means in this context. Indeed, the "best arm" for one agent may not be the best for another. A first attempt may be to associate some "aggregate quality" to each arm; for example, the quality of arm $j$ may be defined as the total mean reward it gives to all agents, i.e., $\sum_i \mu_{i,j}^*$. This would nicely reduce our problem to the classic multi-armed bandit problem, for which we have an armory of available solutions [Slivkins et al., 2019]. However, this approach suffers from the *tyranny of the majority* [Moulin, 2003]. For example, imagine a scenario with ten agents, two arms, and deterministic rewards. Suppose four agents derive a reward of $1$ from the first arm but $0$ from the second, while the remaining six derive a reward of $1$ from the second arm but $0$ from the first. The aforementioned approach will deem the second arm as the best

and a classical MAB algorithm will converge to *repeatedly* pulling the second arm, thus unfairly treating the first four agents (a minority). A solution which treats each group in a "proportionally fair" [Moulin, 2003] manner should ideally converge to pulling the first arm $40\%$ of the time and the second $60\%$ of the time. Alternatively, we can allow the learning algorithm to "pull" a probability distribution over the arms and seek an algorithm that converges to placing probability $0.4$ on the first arm and $0.6$ on the second.

This problem of making a fair collective decision when the available alternatives — in this case, probability distributions over the arms — affect multiple agents is well-studied in computational social choice [Brandt et al., 2016]. The literature offers a compelling fairness notion called the *Nash social welfare*, named after John Nash. According to this criterion, the fairest distribution maximizes the *product* of the expected utilities (rewards) to the agents. A distribution $p$ that places probability $p_j$ on each arm $j$ gives expected utility $\sum_j p_j \cdot \mu^*_{i,j}$ to agent $i$. Hence, the goal is to maximize $\mathrm{NSW}(p, \mu^*) = \prod_{i=1}^{N}(\sum_{j=1}^{K} p_j \cdot \mu^*_{i,j})$ over $p$. One can verify that this approach on the aforementioned example indeed yields probability $0.4$ on the first arm and $0.6$ on the second, as desired. It is also interesting to point out that with a single agent ($N = 1$), the distribution maximizing the Nash social welfare puts probability $1$ on the best arm, thus effectively reducing the problem to the classical multi-armed bandit problem (albeit with subtle differences which we highlight in Section 8.3).

Maximizing the Nash social welfare is often seen as a middle ground between maximizing the utilitarian social welfare (sum of utilities to the agents), which is unfair to minorities (as we observed), and maximizing the egalitarian social welfare (minimum utility to any agent), which is considered too extreme [Moulin, 2003]. The solution maximizing the Nash social welfare is also known to satisfy other qualitative fairness desiderata across a wide variety of settings [Fain et al., 2018, Conitzer et al., 2017, Amanatidis et al., 2020, Caragiannis et al., 2019b, Freeman et al., 2017, Aziz et al., 2019a, Fain et al., 2016, Brandl et al., 2022]. For example, a folklore result shows that in our setting such a solution will always lie in *the core*; we refer the reader to the work of Fain et al. [2018] for a formal definition of the core as well as a short derivation of this fact using the first-order optimality condition. For further discussion on this, see Sections 8.1 and 8.7.

When *exactly* maximizing the Nash social welfare is not possible (either due to a lack of complete information, as in our case, or due to computational difficulty), researchers have sought to achieve approximate fairness by approximately maximizing this objective [Cole et al., 2017, Anari et al., 2018, Lee, 2017, Garg and McGlaughlin, 2019, Cole and Gkatzelis, 2018, Garg et al., 2018]. Following this approach in our problem, we define the (cumulative) regret of an algorithm at horizon $T$ as $\sum_{t=1}^{T}(\max_p \mathrm{NSW}(p, \mu^*) - \mathrm{NSW}(p^t, \mu^*))$, where $p^t$ is the distribution selected in round $t$. Our goal in this chapter is to design algorithms whose regret is sublinear in $T$.

## Our Results

We consider three classic algorithms for the multi-armed bandit problem: Explore-First, Epsilon-Greedy, and UCB [Slivkins et al., 2019].

All three algorithms attempt to balance exploration (pulling arms only to learn their rewards) and exploitation (using the information learned so far to pull "good" arms). Explore-First performs exploration for a number of rounds optimized as a function of $T$ followed by exploitation in the remaining rounds to achieve regret $\widetilde{\mathcal{O}}\left(K^{1/3}T^{2/3}\right)$. Epsilon-Greedy flips a coin in each round to decide whether to perform exploration or exploitation and achieves the same regret bound. Its key advantage over

Explore-First is that it does not need to know the horizon $T$ upfront. UCB merges exploration and exploitation to achieve a regret bound of $\widetilde{\mathcal{O}}\left(K^{1/2}T^{1/2}\right)$. Here, $\widetilde{\mathcal{O}}$ hides log factors. Traditionally, the focus is on optimizing the exponent of $T$ rather than that of $K$ as the horizon $T$ is often much larger than the number of arms $K$. It is known that the dependence of UCB's regret on $T$ is optimal: no algorithm can achieve *instance-independent* $o(T^{1/2})$ regret [Auer et al., 2002].[1]

We propose natural multi-agent variants of these three algorithms. Our variants take the Nash social welfare objective into account and select a distribution over the arms in each round instead of a single arm. For Explore-First, we derive $\widetilde{\mathcal{O}}\left(N^{2/3}K^{1/3}T^{2/3}\right)$ regret bound, which recovers the aforementioned single-agent bound with an additional factor of $N^{2/3}$. We also show that changing a parameter of the algorithm yields a regret bound of $\widetilde{\mathcal{O}}\left(N^{1/3}K^{2/3}T^{2/3}\right)$, which offers a different tradeoff between the dependence on $N$ and $K$. For Epsilon-Greedy, we recover the same two regret bounds, although the analysis becomes much more intricate. This is because, as mentioned above, Epsilon-Greedy is a horizon-independent algorithm (i.e. it does not require apriori knowledge of $T$), unlike Explore-First. For UCB, we derive $\widetilde{\mathcal{O}}\left(NKT^{1/2}\right)$ and $\widetilde{\mathcal{O}}\left(N^{1/2}K^{\frac{3}{2}}T^{1/2}\right)$ regret bounds; our dependence on $K$ worsens compared to the classical single-agent case, but importantly,

we recover the same $\sqrt{T}$ dependence. Finally, we note that even for $N = 1$, a learning algorithm is slightly more powerful in our setting than in the classical setting since it can choose a distribution over the arms as opposed to a deterministic arm. Nonetheless, we can easily derive an $\Omega(\sqrt{KT})$ instance-independent lower bound on the regret of any algorithm in our setting, by adjusting the proof of the same lower bound for the classical multi-armed bandit problem. More details can be found in Section 8.3.

Deriving these regret bounds for the multi-agent case requires overcoming two key difficulties that do not appear in the single-agent case. First, our goal is to optimize a complicated function, the Nash social welfare, rather than simply selecting the best arm. This requires a Lipschitz-continuity analysis of the Nash social welfare function and the use of new tools such as the McDiarmid's inequality which are not needed in the standard analysis. Second, the optimization is over an infinite space (the set of distributions over arms) rather than over a finite space (the set of arms). Thus, certain tricks such as a simple union bound no longer work; we use the concept of $\delta$-covering, used heavily in the Lipschitz bandit framework [Kleinberg et al., 2008], in order to address this.

Our contributions are twofold. Conceptually, we promote a multi-agent viewpoint that requires striking a tradeoff between multiple reward functions; this can be applied to other classical single-agent problems. More technically, as explained below, the existing MAB literature does not provide a $\sqrt{T}$ bound for the Nash social welfare objective. Our regret bounds may be generalizable to a broader family of objectives for which a $\sqrt{T}$ bound was previously unknown.

## Related Work

Since the multi-armed bandit problem was introduced by Thompson [1933], many variants of it have been proposed, such as sleeping bandit [Kleinberg et al., 2010], contextual bandit [Woodroofe, 1979], dueling bandit [Yue et al., 2012], Lipschitz bandit [Kleinberg et al., 2008], etc. However, all these variants involve optimizing the cumulative regret or identifying the Pareto frontier of multiple objectives [Paria

[1] In instance-independent bounds, the constant inside the big-Oh notation is not allowed to depend on the (unknown) distributions in the given instance. UCB also achieves an $O(\log T)$ instance-dependent regret bound, which is also known to be asymptotically optimal [Lai and Robbins, 1985]. For further discussion, see Section 8.7.

et al., 2019] *from the perspective of a single agent*. We note that other multi-agent variants of the multi-armed bandit problem have been explored recently [Chakraborty et al., 2017, Bargiacchi et al., 2018], including in distributed environments [Landgren et al., 2016, Shahrampour et al., 2017, Dubey and Pentland, 2020]. However, they still involve a common reward like in the classical multi-armed bandit problem. Their focus is on getting the agents to cooperate to maximize this common reward.

It is possible to view our work from a single-agent prespective by treating the $K$-simplex of probability distributions over arms as a "continuum of arms" and the Nash welfare as a non-linear objective to be optimized. Kleinberg et al. [2019] explore continuum-armed bandits, but establish a $\widetilde{\mathcal{O}}\left(T^{\frac{\gamma+1}{\gamma+2}}\right)$ regret bound, where $\gamma$ is the zooming dimension, which would be $\Theta(K)$ in our case. Bubeck et al. [2011] also study continuum-armed bandits and establish a similar regret bound, but with respect to the near-optimality dimension, which is again $\Theta(K)$ in our case. Hence, the guarantees of Kleinberg et al. [2019] and Bubeck et al. [2011], when applied to our setting, are worse than our $\widetilde{\mathcal{O}}\left(\sqrt{T}\right)$ regret bound. The literature on bandit convex optimization [Hazan and Levy, 2014, Bubeck and Eldan, 2016, Bubeck et al., 2017] offers $\widetilde{\mathcal{O}}\left(\sqrt{T}\right)$ regret bound, but requires the objective to be concave, which the Nash social welfare (NSW) is not.[2]

Another key aspect of our framework is the focus on fairness. Recently, several papers have focused on fairness in the multi-armed bandit problem. For instance, Joseph et al. [2016] design a UCB variant which guarantees what they refer to as meritocratic fairness to the arms, i.e., that a worse arm is never preferred to a better arm regardless of the algorithm's confidence intervals for them. Liu et al. [2017] require that similar arms be treated similarly, i.e., two arms with similar mean rewards be selected with similar probabilities. Gillen et al. [2018] focus on satisfying fairness with respect to an unknown fairness metric. Finally, Patil et al. [2020] assume that there are external constraints requiring that each arm be pulled in at least a certain fraction of the rounds and design algorithms that achieve low regret subject to this constraint. All these papers seek to achieve fairness *with respect to the arms*. In contrast, in our work, the arms are "inanimate" (e.g. policy decisions) and we seek fairness *with respect to the agents*, who are separate from the arms.

More broadly, the problem of making a fair decision given the (possibly conflicting) preferences of multiple agents is well-studied in computational social choice [Brandt et al., 2016] in a variety of contexts. For example, one can consider our problem as that of randomized voting (alternatively known as *fair mixing* [Aziz et al., 2019a]) by viewing the agents as voters and the arms are candidates. The goal is then to pick a fair lottery over the candidates given the voters' preferences. This is also a special case of other more complex models studied in the literature such as fair public decision-making [Conitzer et al., 2017] and fair allocation of public goods [Fain et al., 2018]. However, in computational social choice, voters typically have fixed preferences over the candidates. In contrast, rewards observed by the agents in our framework are stochastic. From this viewpoint, our work provides algorithms for maximizing the Nash social welfare when noisy information can be queried regarding agent preferences.

[2]Since NSW is log-concave, one might argue that we can apply these results to log-NSW. However, a regret bound in log-NSW does not imply the desired regret bound in NSW.

## 8.2 Preliminaries

For $n \in \mathbb{N}$, define $[n] = \{1, \ldots, n\}$. Let $N, K \in \mathbb{N}$. In the *multi-agent multi-armed bandit* (MA-MAB) problem, there is a set of *agents* $[N]$ and a set of *arms* $[K]$. For each agent $i \in [N]$ and arm $j \in [K]$, there is a *reward distribution* $D_{i,j}$ with mean $\mu^*_{i,j}$ and support $[0, 1]$;[3] when arm $j$ is pulled, each agent $i$ observes an independent *reward* sampled from $D_{i,j}$. Let us refer to $\mu^* = (\mu^*_{i,j})_{i\in[N],j\in[K]} \in [0, 1]^{N\times K}$ as the (true) reward matrix.

**Policies:** As mentioned in the introduction, pulling an arm deterministically may be favorable to one agent, but disastrous to another. Hence, we are interested in *probability distributions* over arms, which we refer to as *policies*. The $K$-simplex, denoted $\Delta^K$, is the set of all policies. For a policy $p \in \Delta^K$, $p_j$ denotes the probability with which arm $j$ is pulled. Note that due to linearity of expectation, the expected reward to agent $i$ under policy $p$ is $\sum_{j=1}^K p_j \cdot \mu^*_{i,j}$.

**Nash social welfare:** The Nash social welfare is defined the product of (expected) rewards to the agents. Given $\mu = (\mu_{i,j})_{i\in[N],j\in[K]}$, and policy $p \in \Delta^K$, define $\mathrm{NSW}(p, \mu) = \prod_{i=1}^N \left(\sum_{j=1}^K p_j \cdot \mu_{i,j}\right)$. Thus, the (true) Nash social welfare under policy $p$ is $\mathrm{NSW}(p, \mu^*)$. Hence, if we knew $\mu^*$, we would pick an *optimal policy* $p^* \in \arg\max_{p\in\Delta^K} \mathrm{NSW}(p, \mu^*)$. However, because we do not know $\mu^*$ in advance, our algorithms will often produce an estimate $\widehat{\mu}$, and use it to choose a policy; the quantity $\mathrm{NSW}(p, \widehat{\mu})$ will play a key role in our algorithms and their analysis.

**Algorithms:** An algorithm for the MA-MAB problem chooses a policy $p^t$ in each round $t \in \mathbb{N}$. Then, an arm $a^t$ is sampled according to policy $p^t$, and for each agent $i \in [N]$, a reward $X^t_{i,a^t}$ is sampled independently from distribution $D_{i,a^t}$. At the end of round $t$, the algorithm learns the sampled arm $a^t$ and the reward vector $(X^t_{i,a^t})_{i\in[N]}$, which it can use to choose policies in the later rounds.

**Reward estimates:** All our algorithms maintain an estimate of the mean reward matrix $\mu^*$ at every round. For round $t$ and arm $j \in [K]$, let $n^t_j = \sum_{s=1}^{t-1} \mathbb{1}[a^s = j]$ denote the number of times arm $j$ is pulled at the beginning of round $t$, and let $\widehat{\mu}^t_{i,j} = \frac{1}{n^t_j} \sum_{s\in[t-1]:a^s=j} X^s_{i,j}$ denote the average reward experienced by agent $i$ from the $n^t_j$ pulls of arm $j$ thus far. Our algorithms treat these as an estimate of $\mu^*_{i,j}$ available at the beginning of round $t$. Let $\widehat{\mu}^t = (\widehat{\mu}^t_{i,j})_{i\in[N],j\in[K]}$.

**Regret:** Recall that $p^*$ is an optimal policy that has the highest Nash social welfare. The *instantaneous regret* in round $t$ due to an algorithm choosing $p^t$ is $r^t = \mathrm{NSW}(p^*, \mu^*) - \mathrm{NSW}(p^t, \mu^*)$. The (cumulative) *regret* in round $T$ due to an algorithm choosing $p^1, \ldots, p^T$ is $R^T = \sum_{t=1}^T r^t$. We note that $R^T$ and $r^t$ are defined for a specific algorithm, which will be clear from the context. We are interested in bounding the *expected regret* $\mathsf{E}[R^T]$ of an algorithm at round $T$, where the expectation is over the randomness involved in sampling the arms $a^t$ and the agent rewards $(X^t_{i,a^t})_{i\in[N]}$ for $t \in [T]$.[4] We say that an algorithm is *horizon-dependent* if it needs to know $T$ in advance, and *horizon-independent* if it yields a regret bound at each possible $T$ simultaneously.

**$\delta$-Covering:** Given a metric space $(X, d)$ and $\delta > 0$, a set $S \subseteq X$ is called a *$\delta$-cover* if for each $x \in X$, there exists $s \in S$ with $d(x, s) \leqslant \delta$. That is, from each point in the metric space, there is a point in the $\delta$-cover that is no more than $\delta$ distance away. We will heavily use the fact that there exists a $\delta$-cover of

[3] We need the support of the distribution to be non-negative and bounded, but the upper bound of 1 is without loss of generality. All our bounds scale linearly with the upper bound on the support.

[4] The algorithms we study do not introduce any further randomness in choosing the policies.

$(\Delta^K, \|\cdot\|_1)$ (i.e. the $K$-simplex under the $L_1$ distance) with size at most $(1 + 2/\delta)^K$ [Wainwright, 2019, p. 126], which follows from a simple discretization of the simplex.

## 8.3 The Lower Bound

In this section, we establish a lower bound on the expected regret of any algorithm for our multi-agent multi-armed bandit (MA-MAB) problem. In the classical multi-armed bandit problem, it is known that no algorithm can achieve a regret bound of $E[R^T] = o(\sqrt{KT})$, when the constant inside the little-Oh notation is required to be independent of the distributions in the given instance [Auer et al., 2002]. For further discussion on bounds where the constant is allowed to depend on the distributions in the given instance, we refer the reader to Section 8.7. Our goal in this section is twofold. First, we argue that, due to subtle reasons, a lower bound for the classical multi-armed bandit problem does not immediately carry over to our setting. Nevertheless, we subsequently argue that a very simple adjustment to the proof of the classical lower bound due to Auer et al. [2002] is sufficient to extend it to our problem. To avoid excessive repetition of notation and proof arguments, we purposefully leave this section *not* self-contained and only outline these adjustments needed. We refer an interested reader to the work of Auer et al. [2002] for the full and detailed proof. That said, this argument establishes that the $\sqrt{T}$-dependence of the expected regret of our UCB variant from Theorem 47 is optimal. Note that our focus is solely on the dependence of the expected regret on $T$ as $T$ is typically much larger than both the number of agents $N$ and the number of arms $K$. We leave it to future work to optimize the dependence on $N$ and $K$.

First, we notice that any lower bound derived for the case of a single agent also holds when there are $N > 1$ agents. This is because one can consider instances in which all but one of the agents derive a fixed reward of $1$ from every arm. Note that the contribution of such agents to the product in the Nash social welfare expression is always $1$ regardless of the policy chosen. Hence, the Nash social welfare reduces to simply the expected utility of the remaining agent, i.e., the Nash social welfare in an instance with only this one agent. Therefore, any lower bound on the expected regret that holds for MA-MAB with a single agent also holds for MA-MAB with $N > 1$ agents.

Next, let us focus on the MA-MAB problem with $N = 1$ agent. At the first glance, this may look almost identical to the classical multi-armed bandit problem. After all, if there is but one agent, the policy maximizing the Nash social welfare places probability $1$ on the arm $j^*$ that gives the highest mean reward to the agent. Thus, like in the classical problem, our goal would be to converge to pulling arm $j^*$ repeatedly and our regret would also be measured with respect to the best policy which deterministically pulls arm $j^*$ in every round. However, there are two subtle differences which prevent us from directly borrowing the classical lower bound.

1. In our MA-MAB problem, an algorithm is allowed to "pull" a *distribution over the arms* $p^t$ in round $t$ and learn the stochastically generated rewards for a *random* arm $j^t$ sampled from this distribution. This makes the algorithm slightly more powerful than an algorithm in the classical MAB problem which must deterministically choose an arm to pull.

2. In our MA-MAB problem, the regret in round $t$ is computed as the difference between the mean reward of the best arm and the *expected* mean reward of an arm $j^t$ sampled according to the

distribution $p^t$ used by the algorithm. In the classical problem, one would replace the latter term with the mean reward of the arm actually pulled in round $t$.

The latter distinction is not particularly troublesome because our focus is on the *expected* regret of an algorithm anyway. However, the first distinction makes it impossible to directly borrow lower bounds from the classical MAB problem.

One might wonder if there is still a way to reduce the MA-MAB problem with $N = 1$ agent to the classical MAB problem. For example, given an algorithm $A$ for MA-MAB with $N = 1$, what if we construct an algorithm $\widehat{A}$ for the classical MAB and use the lower bound on the expected regret of $\widehat{A}$ to derive a lower bound on the expected regret of $A$? The problem with such reduction is that once $A$ chooses a distribution $p^t$, we have no control over which arm will be sampled. This choice is crucial as it will determine what information the algorithm gets to learn. We cannot mimic this learning process in our deterministic algorithm $\widehat{A}$. Upon careful consideration, it also seems difficult to express the expected regret of $A$ as the convex combination of the expected regret of several deterministic algorithms for the classical MAB.

Instead of aiming to find a black-box reduction to the classical problem, we therefore investigate in detail the proof of the $\Omega\left(\sqrt{KT}\right)$ lower bound for the classical MAB due to Auer et al. [2002, Theorem 5.1] and observe that their argument goes through for our MA-MAB problem as well. Instead of repeating their proof, we survey the key steps of their proof in which they assume the algorithm to be deterministically pulling an arm and highlight why the argument holds even when this is not the case.

- In the proof of their Lemma A.1, in the explanation of their Equation (30), they cite the assumption that given the rewards observed in the first $t-1$ rounds (they denote this by the vector $\mathbf{r}^{t-1}$), the algorithm pulls a fixed arm $i_t$ in round $t$. They refer to the distribution $\mathbf{P}_i\{r^t|\mathbf{r}^{t-1}\}$ of the reward in round $t$ given $\mathbf{r}^{t-1}$. In their case, the randomness in $r^t$ is solely due to stochasticity of the rewards since the arm pulled ($i_t$) is fixed. However, in our case, one can think of $\mathbf{P}_i\{r^t|\mathbf{r}^{t-1}\}$ as containing randomness both due to the random choice of $i_t$ and due to the stochasticity of the rewards, and their equations still go through.

- In the same equation, they consider $\mathbf{P}_{unif}\{i_t = i\}$, the probability that arm $i$ is pulled in round $t$. In their case, the only randomness is due to $\mathbf{r}^{t-1}$. In our case, there is additional randomness due to the sampling of an arm in round $t$ from a distribution $p^t$. However, this does not affect their calculations.

- Finally, in the proof of their Theorem A.2, they again consider the probability $\mathbf{P}_i\{i_t = i\}$ and the same argument as above ensures that their proof continues to hold in our setting.

Thus, we have the following lower bound.

**Proposition 8.** *For any algorithm for the MA-MAB problem, there exists a problem instance such that* $\mathsf{E}[R^T] = \Omega\left(\sqrt{KT}\right)$.

## 8.4 Explore-First

**Algorithm 7** Explore-First

**Input:** Number of agents $N$, number of arms $K$, horizon $T$
**Parameters:** Exploration period $L$

**for** $t = 1, \ldots, K \cdot L$ **do** {Exploration (Pull each arm $L$ times)}
  $j \leftarrow \lceil t/L \rceil$
  $p^t \leftarrow$ policy that puts probability 1 on arm $j$ {Pull arm $j$ deterministically}
**end for**
Compute the estimated reward matrix $\widehat{\mu} \triangleq \widehat{\mu}^{K \cdot L+1}$ of the rewards observed so far
Compute $\widehat{p} \in \arg\max_{p \in \Delta^K} \mathrm{NSW}(p, \widehat{\mu})$
**for** $t = K \cdot L + 1, \ldots, T$ **do** {Exploitation}
  $p^t \leftarrow \widehat{p}$
**end for**

Perhaps the simplest algorithm (with a sublinear regret bound) in the classic single-agent MAB framework is Explore-First. It is composed of two distinct stages. The first stage is *exploration*, during which the algorithm pulls each arm $L$ times. At the end of this stage, the algorithm computes the arm $\widehat{a}$ with the best estimated mean reward, and in the subsequent *exploitation* stage, pulls arm $\widehat{a}$ in every round. The algorithm is horizon-dependent, i.e., it takes the horizon $T$ as input and sets $L$ as a function of $T$. Setting $L = \Theta\left(K^{-\frac{2}{3}} T^{\frac{2}{3}} \log^{\frac{1}{3}}(T)\right)$ yields regret bound $\mathsf{E}[R^T] = \mathcal{O}\left(K^{\frac{1}{3}} T^{\frac{2}{3}} \log^{\frac{1}{3}}(T)\right)$ [Slivkins et al., 2019].

In our multi-agent variant, presented as Algorithm 7.

the exploration stage pulls each arm $L$ times as before. However, at the end of this stage, the algorithm computes, not an arm $\widehat{a}$, but a policy $\widehat{p}$ with the best estimated Nash social welfare. During exploitation, it then uses policy $\widehat{p}$ in every round. With an almost identical analysis as in the single-agent setting, we recover the aforementioned regret bound with an additional $N^{2/3}$ factor for $N$ agents.

Using a novel and more intricate argument, we show that a different tradeoff between the exponents of $N$ and $K$ can be obtained, where $N^{2/3}$ is reduced to $N^{1/3}$ at the expense of increasing $K^{1/3}$ to $K^{2/3}$ (and adding a logarithmic term). We later use this approach again in our analysis of more sophisticated algorithms.

We remark that Algorithm 7 can be implemented efficiently. The only non-trivial step is to compute the optimal policy given an estimated reward matrix, i.e., $\widehat{p} \in \arg\max_{p \in \Delta^K} \mathrm{NSW}(p, \widehat{\mu})$. Since the Nash social welfare is known to be log-concave, this can be solved in polynomial time [Eisenberg and Gale, 1959].

Our analysis of Explore-First and later algorithms relies on a few elementary lemmas regarding the behavior of the Nash social welfare function $\mathrm{NSW}(p, \mu)$. We are mainly interested in how much the function can change when its arguments change. To that end, the following folklore result translates the difference in a product to a sum of point-wise differences that are easier to deal with.

**Lemma 21.** *Let* $a_i, b_i \in [0, 1]$ *for* $i \in [N]$. *Then,* $\left|\prod_{i=1}^N a_i - \prod_{i=1}^N b_i\right| \leqslant \sum_{i=1}^N |a_i - b_i|$.

*Proof.* We prove this using induction on $N$. For $N = 1$, the lemma trivially holds. Suppose it holds for

$N = n$. For $N = n+1$, we have

$$\begin{aligned}\left|\prod_{i=1}^{n+1} a_i - \prod_{i=1}^{n+1} b_i\right| &= \left|\prod_{i=1}^{n+1} a_i - b_{n+1}\prod_{i=1}^{n} a_i + b_{n+1}\prod_{i=1}^{n} a_i - \prod_{i=1}^{n+1} b_i\right| \\ &\leqslant \left(\prod_{i=1}^{n} a_i\right)|a_{n+1} - b_{n+1}| + b_{n+1}\cdot\left|\prod_{i=1}^{n} a_i - \prod_{i=1}^{n} b_i\right| \\ &\leqslant |a_{n+1} - b_{n+1}| + \sum_{i=1}^{n}|a_i - b_i| = \sum_{i=1}^{n+1}|a_i - b_i|\,,\end{aligned}$$

where the second transition is due to the triangle inequality, and the third transition holds due to the induction hypothesis and because $a_i, b_i \in [0,1]$ for each $i$. □

Using Lemma 21, we can easily analyze Lipschitz-continuity of $\mathrm{NSW}(p,\mu)$ when either $p$ or $\mu$ changes and the other is fixed. First, we consider change in $p$ with $\mu$ fixed.

**Lemma 22.** *Given a reward matrix $\mu \in [0,1]^{N\times K}$ and policies $p^1, p^2 \in \Delta^K$, we have*

$$\left|\mathrm{NSW}(p^1,\mu) - \mathrm{NSW}(p^2,\mu)\right| \leqslant N\cdot\left\|p^1 - p^2\right\|_1 = N\cdot\textstyle\sum_{j\in[K]}\left|p_j^1 - p_j^2\right|.$$

*Proof.* Using Lemma 21, we have

$$|\mathrm{NSW}(p^1,\mu) - \mathrm{NSW}(p^2,\mu)| \leqslant \textstyle\sum_{i\in[N]}\left|\sum_{j\in[K]}(p_j^1 - p_j^2)\cdot\mu_{i,j}\right| \leqslant N\cdot\sum_{j\in[K]}\left|p_j^1 - p_j^2\right|,$$

where the final transition is due to the triangle inequality and because $\mu_{i,j} \in [0,1]$ for each $i, j$. □

Next, we consider change in $\mu$ with $p$ fixed.

**Lemma 23.** *Given a policy $p \in \Delta^K$, and reward matrices $\mu^1, \mu^2 \in [0,1]^{N\times K}$, we have*

$$\left|\mathrm{NSW}(p,\mu^1) - \mathrm{NSW}(p,\mu^2)\right| \leqslant \textstyle\sum_{i\in[N]}\sum_{j\in[K]} p_j\cdot\left|\mu_{i,j}^1 - \mu_{i,j}^2\right|.$$

*Proof.* Again, using Lemma 21, we have

$$\begin{aligned}&|\mathrm{NSW}(p,\mu^1) - \mathrm{NSW}(p,\mu^2)| \\ &\leqslant \sum_{i\in[N]}\left|\sum_{j\in[K]} p_j\cdot(\mu_{i,j}^1 - \mu_{i,j}^2)\right| \leqslant \sum_{i\in[N],j\in[K]} p_j\cdot\left|\mu_{i,j}^1 - \mu_{i,j}^2\right|,\end{aligned}$$

where the last transition is due to the triangle inequality. □

We are now ready to establish the regret bounds for Explore-First.

**Theorem 45.** *Explore-First is horizon-dependent and has the following expected regret at round $T$.*

- *When $L = \Theta\left(N^{\frac{2}{3}}K^{-\frac{2}{3}}T^{\frac{2}{3}}\log^{\frac{1}{3}}(NKT)\right)$, $\mathsf{E}[R^T] = \mathcal{O}\left(N^{\frac{2}{3}}K^{\frac{1}{3}}T^{\frac{2}{3}}\log^{\frac{1}{3}}(NKT)\right)$.*

- *When* $L = \Theta\left(N^{\frac{1}{3}}K^{-\frac{1}{3}}T^{\frac{2}{3}}\log^{\frac{2}{3}}(NKT)\right)$, $\mathsf{E}[R^T] = \mathcal{O}\left(N^{\frac{1}{3}}K^{\frac{2}{3}}T^{\frac{2}{3}}\log^{\frac{2}{3}}(NKT)\right)$.

*Proof.* Note that the instantaneous regret $r^t(p^t)$ in any round $t$ can be at most 1 because $\mathrm{NSW}(p, \mu^*) \in [0, 1]$ for every policy $p$. Thus,

$$\mathsf{E}[R^T] = \sum_{t=1}^{T} \mathsf{E}[r^t] \leqslant KL \cdot 1 + (T - KL) \cdot \mathsf{E}[\mathrm{NSW}(p^*, \mu^*) - \mathrm{NSW}(\widehat{p}, \mu^*)]. \tag{8.1}$$

Thus, our goal is to bound $\mathsf{E}[\mathrm{NSW}(p^*, \mu^*) - \mathrm{NSW}(\widehat{p}, \mu^*)]$. We bound this in two ways.

**First approach:** We present this approach briefly since it largely mimics the classical analysis with an application of Lemma 23. Here, we bound how much $\widehat{\mu}$ can deviate from $\mu^*$. Specifically, we let $\epsilon = \sqrt{\frac{\log(NKT)}{L}}$ and define the event $\mathcal{E} \triangleq \forall i \in [N], \forall j \in [K] : \left|\widehat{\mu}_{i,j} - \mu^*_{i,j}\right| \leqslant \epsilon$. Since $L$ is fixed, we have $\mathsf{E}[\widehat{\mu}_{i,j}] = \mu^*_{i,j}$. Hence, we can apply Hoeffding's inequality followed by the union bound to derive $\Pr[\mathcal{E}] \geqslant 1 - 2/T^2$. Conditioned on $\mathcal{E}$, from Lemma 23 we have $\mathrm{NSW}(p, \mu^*) - \mathrm{NSW}(p, \widehat{\mu}) \leqslant N\epsilon$ for every policy $p$, which implies

$$\mathrm{NSW}(p^*, \mu^*) \leqslant \mathrm{NSW}(p^*, \widehat{\mu}) + N\epsilon \leqslant \mathrm{NSW}(\widehat{p}, \widehat{\mu}) + N\epsilon \leqslant \mathrm{NSW}(\widehat{p}, \mu^*) + 2N\epsilon,$$

where the second transition is because $\widehat{p} \in \arg\max_{p \in \Delta^K} \mathrm{NSW}(p, \widehat{\mu})$. Substituting this into Equation (8.1), using the fact that $\mathsf{E}[\mathrm{NSW}(p^*, \mu^*) - \mathrm{NSW}(\widehat{p}, \mu^*)] \leqslant 1 \cdot \mathsf{E}[\mathrm{NSW}(p^*, \mu^*) - \mathrm{NSW}(\widehat{p}, \mu^*)|\mathcal{E}] + \Pr[\neg\mathcal{E}] \cdot 1$, and setting $L = \Theta\left(N^{\frac{2}{3}}K^{-\frac{2}{3}}T^{\frac{2}{3}}\log^{\frac{1}{3}}(NKT)\right)$ yields the first regret bound.

**Second approach:** We now focus on another approach for bounding $\mathsf{E}[\mathrm{NSW}(p^*, \mu^*) - \mathrm{NSW}(\widehat{p}, \mu^*)]$, which is more intricate and offers a different tradeoff between the dependence on $N$ and $K$. Notice that for a given $p$, $\mathsf{E}[\mathrm{NSW}(p, \widehat{\mu})] = \mathrm{NSW}(p, \mu^*)$ because all $\widehat{\mu}_{i,j}$-s are independent and expectation decomposes over sums and products of independent random variables. Thus, we can use McDiarmid's inequality to bound $|\mathrm{NSW}(p, \widehat{\mu}) - \mathrm{NSW}(p, \mu^*)|$ at a given $p$.

Fix a $\delta$-cover $\mathcal{P}$ of $(\Delta^K, \|\cdot\|_1)$ with $|\mathcal{P}| \leqslant (1 + 2/\delta)^K$. Fix $p \in \mathcal{P}$. Notice that $\widehat{\mu}_{i,j} = (1/L) \cdot \sum_{s=1}^{L} X^s_{i,j}$, where $X^s_{i,j}$ is the reward to agent $i$ from the $s$-th pull of arm $j$ during the exploration phase.

We thus decompose $\widehat{\mu}$ into $N \cdot L$ random variables: for each $i \in [N]$ and $s \in [L]$, we let $X^s_i = (X^s_{i,j})_{j \in [K]}$. To apply McDiarmid's inequality, we need to analyze the maximum amount $c^s_i$ by which changing $X^s_i$ can change $\mathrm{NSW}(p, \widehat{\mu})$. Using Lemma 23, it is easy to see that $c^s_i \leqslant 1/L$ for each $i \in [N]$ and $s \in [L]$. Now, applying McDiarmid's inequality, we have

$$\Pr\left[|\mathrm{NSW}(p, \widehat{\mu}) - \mathrm{NSW}(p, \mu^*)| \leqslant \epsilon\right] \leqslant 2e^{\frac{-2\epsilon^2}{\sum_{i \in [N], s \in [L]} (c^s_i)^2}} = 2e^{\frac{-2L\epsilon^2}{N}}.$$

Setting $\epsilon = \sqrt{\frac{N \log(|\mathcal{P}|T)}{2L}}$, we have that for each $p \in \mathcal{P}$,

$$\Pr\left[|\mathrm{NSW}(p, \widehat{\mu}) - \mathrm{NSW}(p, \mu^*)| \leqslant \sqrt{\frac{N \log(|\mathcal{P}|T)}{2L}}\right] \leqslant \frac{2}{|\mathcal{P}|T}.$$

Using the union bound, we have that

$$\Pr\left[\forall p \in \mathcal{P} : |\mathrm{NSW}(p, \widehat{\mu}) - \mathrm{NSW}(p, \mu^*)| \leqslant \sqrt{\frac{N \log(|\mathcal{P}|T)}{2L}}\right] \geqslant 1 - \frac{2}{T}.$$

For $p \in \Delta^K$, let $\overline{p} \in \arg\min_{p' \in \mathcal{P}} \|p - p'\|_1$. Then, since $\mathcal{P}$ is a $\delta$-cover, we have $\|p - \overline{p}\|_1 \leqslant \delta$. Thus, due to Lemma 22, we have

$$\begin{aligned}|\mathrm{NSW}(p, \widehat{\mu}) - \mathrm{NSW}(p, \mu^*)| &\leqslant \sum_{\mu \in \{\widehat{\mu}, \mu^*\}} |\mathrm{NSW}(p, \mu) - \mathrm{NSW}(\overline{p}, \mu)| + |\mathrm{NSW}(\overline{p}, \widehat{\mu}) - \mathrm{NSW}(\overline{p}, \mu^*)| \\ &\leqslant 2N\delta + |\mathrm{NSW}(\overline{p}, \widehat{\mu}) - \mathrm{NSW}(\overline{p}, \mu^*)|\,.\end{aligned}$$

Setting $\delta = \frac{1}{NT}$, we have

$$\Pr\left[\forall p \in \Delta^K : |\mathrm{NSW}(p, \widehat{\mu}) - \mathrm{NSW}(p, \mu^*)| \leqslant \frac{2}{T} + \sqrt{\frac{N \log(|\mathcal{P}|T)}{2L}}\right] \geqslant 1 - \frac{2}{T}.$$

Next, we use the fact that

$$\mathrm{NSW}(p^*, \mu^*) - \mathrm{NSW}(\widehat{p}, \mu^*) \leqslant \sum_{p \in \{p^*, \widehat{p}\}} |\mathrm{NSW}(p, \widehat{\mu}) - \mathrm{NSW}(p, \mu^*)|\,.$$

Hence,

$$\Pr\left[|\mathrm{NSW}(p^*, \mu^*) - \mathrm{NSW}(\widehat{p}, \mu^*)| \leqslant \frac{4}{T} + \sqrt{\frac{2N \log(|\mathcal{P}|T)}{L}}\right] \geqslant 1 - \frac{2}{T}.$$

Next, we substitute $|\mathcal{P}| \leqslant (1 + 2/\delta)^K \leqslant (3/\delta)^K$, $\delta = \frac{1}{NT}$, and $L = \Theta\left(N^{\frac{1}{3}} K^{-\frac{1}{3}} T^{\frac{2}{3}} \log^{\frac{2}{3}}(NKT)\right)$, and then substitute the derived bound in Equation (8.1) to get the second regret bound. □

## 8.5 Epsilon-Greedy

**Algorithm 8** $\epsilon^t$-Greedy

**Input:** Number of agents $N$, number of arms $K$
**Parameters:** Exploration probabilities $\epsilon^t$ for $t \in \mathbb{N}$

$curr \leftarrow 1$ {Next arm to pull during exploration}
**for** $t = 1, 2, \ldots,$ **do**
  Toss a coin with success probability $\epsilon^t$
  **if** success **then** {Exploration (Round-robin among arms)}
    $p^t \leftarrow$ policy that puts probability 1 on arm $curr$ {Pull it deterministically}
    $curr \leftarrow curr + 1$ {When $curr$ becomes $K + 1$, reset to 1}
  **else** {Exploitation}
    Compute the estimated reward matrix $\widehat{\mu}^t$ from the rewards observed so far
    $p^t \leftarrow \arg\max_{p \in \Delta^K} \mathrm{NSW}(p, \widehat{\mu}^t)$
  **end if**
**end for**

A slightly more sophisticated algorithm than Explore-First is Epsilon-Greedy, which is presented as Algorithm 8.

It spreads out exploration instead of performing it all at the beginning. Specifically, at each round $t$, it performs exploration with probability $\epsilon^t$, and exploitation otherwise. Exploration cycles through the arms in a round-robin fashion, while exploitation uses the policy $p^t$ with the highest Nash social welfare under the current estimated reward matrix (rather than choosing a single estimated best arm as in the classical algorithm).

Like Explore-First, Epsilon-Greedy can be implemented efficiently. The only non-trivial step is to compute $\widehat{p} \in \arg\max_{p \in \Delta^K} \mathrm{NSW}(p, \widehat{\mu})$, which, as mentioned before, can be done efficiently.

The key advantage of Epsilon-Greedy over Explore-First is that it is horizon-independent. However, in the $\widehat{\mu}$ computed in Explore-First at the end of exploration, each $\widehat{\mu}_{i,j}$ is the average of $L$ iid samples, where $L$ is fixed. In contrast, in the $\widehat{\mu}^t$ computed in Epsilon-Greedy in round $t$, each $\widehat{\mu}^t_{i,j}$ is the average of $n^t_j$ iid samples. The fact that $n^t_j$ is itself a random variable and the $\widehat{\mu}^t_{i,j}$-s are correlated through the $n^t_j$-s prevents a direct application of certain statistical inequalities, thus complicating the analysis of Epsilon-Greedy. To address this, we first present a sequence of useful lemmas that apply to *any* algorithm, and then use them to prove the regret bounds of Epsilon-Greedy and later UCB.

## Useful Lemmas

Recall that $\mu^*$ and $\widehat{\mu}^t$ denote the true reward matrix and the estimated reward matrix at the beginning of round $t$, respectively. Our goal is to find an upper bound on the quantity $|\mathrm{NSW}(p, \mu^*) - \mathrm{NSW}(p, \widehat{\mu}^t)|$ that, with high probability, holds at every $p \in \Delta^K$ simultaneously. To that end, we first need to show that $\widehat{\mu}^t$ will be close to $\mu^*$ with high probability.

Recall that random variable $n^t_j$ denotes the number of times arm $j$ is pulled by an algorithm before round $t$, and $\widehat{\mu}^t_{i,j}$ is an average over $n^t_j$ independent samples. Hence, we cannot directly apply Hoeffding's inequality, but we can nonetheless use standard tricks from the literature.

**Lemma 24.** *Define* $r^t_j = \sqrt{\frac{2\log(NKt)}{n^t_j}}$*, and event*

$$\mathcal{E}^t \triangleq \forall i \in [N], j \in [K] : \left|\widehat{\mu}^t_{i,j} - \mu^*_{i,j}\right| \leqslant r^t_j.$$

*Then, for any algorithm and any* $t$*, we have* $\Pr[\mathcal{E}^t] \geqslant 1 - \frac{2}{t^3}$.

*Proof.* Fix $t$. For $i \in [N]$, $j \in [K]$, and $\ell \in [t]$, let $\overline{v}^\ell_{i,j}$ denote the average reward to agent $i$ from the first $\ell$ pulls of arm $j$, and define $\overline{r}^\ell_j = \sqrt{\frac{2\log(NKt)}{\ell}}$. Then, by Hoeffding's inequality, we have

$$\forall i \in [N], j \in [K], \ell \in [t] : \Pr\left[\left|\overline{v}^\ell_{i,j} - \mu_{i,j}\right| > \overline{r}^\ell_j\right] \leqslant \frac{2}{(NKt)^4}.$$

By the union bound, we get

$$\Pr\left[\forall i \in [N], j \in [K], \ell \in [t] : \left|\overline{v}^\ell_{i,j} - \mu_{i,j}\right| \leqslant \overline{r}^\ell_j\right] \geqslant 1 - \frac{2}{(NKt)^3}.$$

Because $n^t_j \in [t]$ for each $j \in [K]$, the above event implies our desired event $\mathcal{E}^t$. Hence, we have that $\Pr[\mathcal{E}^t] \geqslant 1 - 2/(NKt)^3 \geqslant 1 - 2/t^3$. □

Conditioned on $\mathcal{E}^t$, we wish to bound $|\mathrm{NSW}(p,\mu^*) - \mathrm{NSW}(p,\widehat{\mu}^t)|$ simultaneously at all $p \in \Delta^K$. We provide two such (incomparable) bounds, which will form the crux of our regret bound analysis. The first bound is a direct application of the Lipschitz-continuity analysis from Lemma 23.

**Lemma 25.** *Conditioned on $\mathcal{E}^t$, we have that*

$$\forall p \in \Delta^K : \left|\mathrm{NSW}(p,\widehat{\mu}^t) - \mathrm{NSW}(p,\mu^*)\right| \leqslant N \cdot \textstyle\sum_{j\in[K]} p_j \cdot r_j^t.$$

*Proof.* Conditioned on $\mathcal{E}^t$, we have $\left|\widehat{\mu}_{i,j}^t - \mu_{i,j}^*\right| \leqslant r_j^t$ for each $j \in [K]$. In that case, it is easy to see that the upper bound from Lemma 23 becomes $N \cdot \sum_{j\in[K]} p_j \cdot r_j^t$. □

The factor of $N$ in Lemma 25 stems from analyzing how much $\widehat{\mu}^t$ may deviate from $\mu^*$ conditioned on $\mathcal{E}^t$, in the worst case. However, even after conditioning on $\mathcal{E}^t$, $\widehat{\mu}^t$ remains a random variable. Hence, one may expect that its deviation, and thus the difference $|\mathrm{NSW}(p,\widehat{\mu}^t) - \mathrm{NSW}(p,\mu^*)|$, may be smaller in expectation. Thus, to derive a different bound than in Lemma 25, we wish to apply McDiarmid's inequality. However, there are two issues in doing so directly.

- McDiarmid's inequality bounds the deviation of $\mathrm{NSW}(p,\widehat{\mu}^t)$ from its expected value. If $\widehat{\mu}^t$ consisted of independent random variables, like in Explore-First, this would be equal to $\mathrm{NSW}(p,\mu^*)$. However, in general, these variables may be correlated through $n_j^t$. We use a conditioning trick to address this issue.
- We cannot hope to apply McDiarmid's inequality at each $p \in \Delta^K$ separately and use the union bound because $\Delta^K$ is infinite. So we apply it at each $p$ in a $\delta$-cover of $\Delta^K$, apply the union bound, and then translate the guarantee to nearby $p \in \Delta^K$ using the Lipschitz-continuity analysis from Lemma 22.

The next result is one of the key technical contributions of our work with a rather long proof.

**Lemma 26.** *Define the event*

$$\mathcal{H}^t \triangleq \forall p \in \Delta^K : \left|\mathrm{NSW}(p,\widehat{\mu}^t) - \mathrm{NSW}(p,\mu^*)\right| \leqslant \sqrt{12NK\log(NKt)} \cdot \textstyle\sum_{j\in[K]} p_j \cdot r_j^t + \dfrac{4}{t}.$$

*Then, for any algorithm and any $t$, we have $\Pr[\mathcal{H}^t|\mathcal{E}^t] \geqslant 1 - 2/t^3$.*

*Proof.* Fix $p \in \Delta^K$. Fix $\delta > 0$, and let $\mathcal{P}$ be a $\delta$-cover of the policy simplex $\Delta^K$ with $|\mathcal{P}| \leqslant (1+2/\delta)^K$ [Wainwright, 2019, p. 126].

Conditioned on $\mathcal{E}^t$ (i.e. $\left|\widehat{\mu}_{i,j}^t - \mu_{i,j}^*\right| \leqslant r_j^t = \sqrt{\frac{2\log(NKt)}{n_j^t}}, \forall i \in [N], j \in [K]$), we wish to derive a high probability bound on $|\mathrm{NSW}(p,\widehat{\mu}^t) - \mathrm{NSW}(p,\mu^*)|$. We can bound the deviation of $\mathrm{NSW}(p,\widehat{\mu}^t)$ from its expected value. However, unlike in the case of Explore-First, we cannot directly claim that the expected value is $\mathrm{NSW}(p,\mu^*)$ because, as we mentioned above, $\widehat{\mu}^t$ consists of random variables that may be correlated through the random varaible $n^t = (n_1^t, \ldots, n_K^t)$ taking values in $[t]^K$. Thus, we need a more careful argument.

For $i \in [N]$, $j \in [K]$, and $\ell_j \in [t]$, let $\overline{v}_{i,j}^{\ell_j}$ denote the average reward to agent $i$ from the first $\ell_j$ pulls of arm $j$, and define $\overline{r}_j^{\ell_j} = \sqrt{\frac{2\log(NKt)}{\ell_j}}$. Let $\ell = (\ell_1, \ldots, \ell_K) \in [t]^K$ and $\overline{v}^\ell = (\overline{v}_{i,j}^{\ell_j})_{i\in[N],j\in[K]}$. Each $\overline{v}_{i,j}^{\ell_j}$

is independent and satisfies $\mathsf{E}[\overline{v}_{i,j}^{\ell_j}] = \mu^*_{i,j}$. Since expectation decomposes over sums and products of independent random variables, we have $\mathsf{E}[\mathrm{NSW}(p, \overline{v}^\ell)] = \mathrm{NSW}(p, \mu^*)$.

**Evaluating conditional expectation:** We next argue that further conditioning on the high probability event $\mathcal{E}^t$ does not change the expectation by much. Formally,

$$\begin{aligned}
&\left|\mathrm{NSW}(p, \mu^*) - \mathsf{E}[\mathrm{NSW}(p, \overline{v}^\ell)|\mathcal{E}^t]\right| \\
&= \left|\mathsf{E}[\mathrm{NSW}(p, \overline{v}^\ell)] - \mathsf{E}\left[\mathrm{NSW}(p, \overline{v}^\ell)|\mathcal{E}^t\right]\right| \\
&= \Pr[\neg\mathcal{E}^t] \cdot \left|\mathsf{E}\left[\mathrm{NSW}(p, \overline{v}^\ell)|\neg\mathcal{E}^t\right] - \mathsf{E}\left[\mathrm{NSW}(p, \overline{v}^\ell)|\mathcal{E}^t\right]\right| \\
&\leqslant \Pr[\neg\mathcal{E}^t] \leqslant \frac{2}{t^3} \leqslant \frac{2}{t},
\end{aligned} \tag{8.2}$$

where the penultimate transition holds because NSW is bounded in $[0, 1]$, and the final transition is due to Lemma 24.

**Applying McDiarmid's inequality:** We first decompose $\overline{v}^\ell$ into $N$ random variables: for each $i \in [N]$, let $\overline{v}_i^\ell = (\overline{v}_{i,j}^\ell)_{j\in[K]}$. To apply McDiarmid's inequality, we need to analyze the maximum amount $c_i$ by which changing $\overline{v}_i^\ell$ can change $\mathrm{NSW}(p, \overline{v}^\ell)$. Fix $i \in [N]$, and fix all the variables except $\overline{v}_i^\ell$. Conditioned on $\mathcal{E}^t$, each $\overline{v}_{i,j}^\ell$ can change by at most $2\overline{r}_j^{\ell_j}$. Hence, using Lemma 23, we have that $c_i \leqslant 2\sum_{j\in[K]} p_j \cdot \overline{r}_j^{\ell_j}$. Now, applying McDiarmid's inequality, we have $\forall \ell \in [t]^K$ :

$$\Pr\left[\left|\mathrm{NSW}(p, \overline{v}^\ell) - \mathsf{E}\left[\mathrm{NSW}(p, \overline{v}^\ell)|\mathcal{E}^t\right]\right| \geqslant \epsilon \mid \mathcal{E}^t\right] \leqslant 2e^{\frac{-2\epsilon^2}{\sum_{i\in[N]} c_i^2}} \leqslant 2e^{\frac{-2\epsilon^2}{4N\cdot\left(\sum_{j\in[K]} p_j \cdot \overline{r}_j^{\ell_j}\right)^2}}.$$

Using Equation (8.2), and setting $\epsilon = \sqrt{2N\log(|\mathcal{P}|t^{K+3})} \cdot \sum_{j\in[K]} p_j \cdot \overline{r}_j^{\ell_j}$, we have that $\forall \ell \in [t]^K$:

$$\Pr\left[\left|\mathrm{NSW}(p, \overline{v}^\ell) - \mathrm{NSW}(p, \mu^*)\right| \geqslant \sqrt{2N\log(|\mathcal{P}|t^{K+3})} \cdot \sum_{j\in[K]} p_j \cdot \overline{r}_j^{\ell_j} + \frac{2}{t} \;\middle|\; \mathcal{E}^t\right] \leqslant \frac{2}{|\mathcal{P}|t^{K+3}}.$$

Next, by union bound, we get

$$\begin{aligned}
&\Pr\left[\forall \ell \in [t]^K : \left|\mathrm{NSW}(p, \overline{v}^\ell) - \mathrm{NSW}(p, \mu^*)\right| \geqslant \sqrt{2N\log(|\mathcal{P}|t^{K+3})} \cdot \sum_{j\in[K]} p_j \cdot \overline{r}_j^{\ell_j} + \frac{2}{t} \;\middle|\; \mathcal{E}^t\right] \\
&\leqslant \frac{2}{|\mathcal{P}|t^3}.
\end{aligned}$$

Because $n_j^t \in [t]$ for each $j \in [K]$, we have

$$\begin{aligned}
&\Pr\left[\left|\mathrm{NSW}(p, \widehat{\mu}^t) - \mathrm{NSW}(p, \mu^*)\right| \geqslant \sqrt{2N\log(|\mathcal{P}|t^{K+3})} \cdot \sum_{j\in[K]} p_j \cdot r_j^t + \frac{2}{t} \;\middle|\; \mathcal{E}^t\right] \\
&\leqslant \frac{2}{|\mathcal{P}|t^3}.
\end{aligned}$$

**Extending to all policies in $\mathcal{P}$:** Using the union bound, we have that

$$\Pr\left[\forall p \in \mathcal{P} : \left|\mathrm{NSW}(p, \widehat{\mu}^t) - \mathrm{NSW}(p, \mu^*)\right| \leqslant \sqrt{2N \log(|\mathcal{P}| t^{K+3})} \cdot \sum_{j \in [K]} p_j \cdot r_j^t + \frac{2}{t} \,\middle|\, \mathcal{E}^t\right] \geqslant 1 - \frac{2}{t^3}.$$

**Extending to all policies in $\mathbf{\Delta^K}$:** For $p \in \Delta^K$, let $\overline{p} \in \arg\min_{p' \in \mathcal{P}} \|p - p'\|_1$. Then, since $\mathcal{P}$ is a $\delta$-cover, we have $\|p - \overline{p}\|_1 \leqslant \delta$. Thus, due to Lemma 22, we have

$$\begin{aligned}\left|\mathrm{NSW}(p, \widehat{\mu}^t) - \mathrm{NSW}(p, \mu^*)\right| &\leqslant \sum_{\mu \in \{\widehat{\mu}^t, \mu^*\}} \left|\mathrm{NSW}(p, \mu) - \mathrm{NSW}(\overline{p}, \mu)\right| \\ &\quad + \left|\mathrm{NSW}(\overline{p}, \widehat{\mu}^t) - \mathrm{NSW}(\overline{p}, \mu^*)\right| \\ &\leqslant 2N\delta + \left|\mathrm{NSW}(\overline{p}, \widehat{\mu}^t) - \mathrm{NSW}(\overline{p}, \mu^*)\right|.\end{aligned}$$

Setting $\delta = \frac{1}{Nt}$, we have

$$\Pr\left[\forall p \in \Delta^K : \left|\mathrm{NSW}(p, \widehat{\mu}^t) - \mathrm{NSW}(p, \mu^*)\right| \leqslant \sqrt{2N \log(|\mathcal{P}| t^{K+3})} \cdot \sum_{j \in [K]} p_j \cdot r_j^t + \frac{4}{t} \,\middle|\, \mathcal{E}^t\right] \geqslant 1 - \frac{2}{t^3}.$$

Substituting $|\mathcal{P}| \leqslant (1 + 2/\delta)^K \leqslant (3/\delta)^K$ with $\delta = \frac{1}{Nt}$ yields the desired bound. □

Finally, we use the following simple observation in deriving our asymptotic bounds.

**Proposition 9.** *For constant $p \in \mathbb{R}$, $\sum_{t=1}^T t^p$ is $\Theta(\log T)$ if $p = -1$ and $\Theta(T^{p+1})$ when $p > -1$.*

## Analysis of Epsilon-Greedy

We can now use these lemmas to establish the regret bounds for Epsilon-Greedy.

**Theorem 46.** *Epsilon-Greedy is horizon-independent, and has the following expected regret at any round $T$.*

- *If $\epsilon^t = \Theta\left(N^{\frac{2}{3}} K^{\frac{1}{3}} t^{-\frac{1}{3}} \log^{\frac{1}{3}}(NKt)\right)$ for all $t$, $\mathsf{E}[R^T] = \mathcal{O}\left(N^{\frac{2}{3}} K^{\frac{1}{3}} T^{\frac{2}{3}} \log^{\frac{1}{3}}(NKT)\right)$.*
- *If $\epsilon^t = \Theta\left(N^{\frac{1}{3}} K^{\frac{2}{3}} t^{-\frac{1}{3}} \log^{\frac{2}{3}}(NKt)\right)$ for all $t$, $\mathsf{E}[R^T] = \mathcal{O}\left(N^{\frac{1}{3}} K^{\frac{2}{3}} T^{\frac{2}{3}} \log^{\frac{2}{3}}(NKT)\right)$.*

*Proof.* Fix $t \in [T]$. Let $b^t$ denote the number of times Epsilon-Greedy performs exploration up to round $t$. Note that $\mathsf{E}[b^t] = \sum_{s=1}^t \epsilon^s \geqslant t\epsilon^t$, where the last step follows from the fact that $\epsilon^t$ is monotonically decreasing in both cases of the theorem. Let $\theta > 0$ be a constant such that $\epsilon^t \geqslant \theta \cdot t^{-1/3}$ in both cases of the theorem.

Define the event $\mathcal{B}^t \triangleq b^t \geqslant \gamma \cdot t\epsilon^t$, where $\gamma = 1 - 1/\theta$. Then, by Hoeffding's inequality, we have

$$\Pr[\neg\mathcal{B}^t] \leqslant e^{-2(1-\gamma)^2\theta^2 t^{1/3}} = e^{-2t^{1/3}} \leqslant e^{-\log t} = \frac{1}{t}. \tag{8.3}$$

Because the algorithm performs round-robin during exploration, conditioned on $\mathcal{B}^t$, we have that $n_j^t \geqslant \frac{b^t}{K} \geqslant \frac{\gamma \cdot t\epsilon^t}{K}$ for each arm $j$,[5] which implies $r_j^t \leqslant \sqrt{\frac{2K\log(NKt)}{\gamma \cdot t\epsilon^t}}$ for each $j$. Thus, conditioned on $\mathcal{B}^t$, we have

$$\forall p \in \Delta^K : \sum_{j\in[K]} p_j \cdot r_j^t \leqslant \max_{j\in[K]} r_j^t \leqslant \sqrt{\frac{2K\log(NKt)}{\gamma \cdot t\epsilon^t}}. \tag{8.4}$$

We are now ready to use the bounds from Lemmas 25 and 26. We focus on the event

$$\mathcal{C}_\alpha^t \triangleq \forall p \in \Delta^K : \left|\mathrm{NSW}(p,\mu^*) - \mathrm{NSW}(p,\widehat{\mu}^t)\right| \leqslant \alpha^t \cdot \sum_{j\in[K]} p_j \cdot r_j^t + \frac{4}{t}.$$

Conditioned on $\mathcal{E}^t \wedge \mathcal{H}^t$, note that $\mathcal{C}_\alpha^t$ holds for $\alpha^t = N$ due to Lemma 25, and for $\alpha^t = \sqrt{12NK\log NKt}$ due to Lemma 26.

Let $\widehat{p}^t \in \arg\max_{p\in\Delta^K} \mathrm{NSW}(p,\widehat{\mu}^t)$. We wish to bound the regret $\mathrm{NSW}(p^*,\mu^*) - \mathrm{NSW}(\widehat{p}^t,\mu^*)$ that Epsilon-Greedy incurs when performing exploitation in round $t$ by choosing policy $\widehat{p}^t$. Conditioned on $\mathcal{E}^t \wedge \mathcal{H}^t \wedge \mathcal{B}^t$, we have

$$\begin{aligned}
&\mathrm{NSW}(p^*,\mu^*) - \mathrm{NSW}(\widehat{p}^t,\mu^*) \\
&= \big(\mathrm{NSW}(p^*,\mu^*) - \mathrm{NSW}(p^*,\widehat{\mu}^t)\big) + \big(\mathrm{NSW}(p^*,\widehat{\mu}^t) - \mathrm{NSW}(\widehat{p}^t,\widehat{\mu}^t)\big) \\
&\qquad + \big(\mathrm{NSW}(\widehat{p}^t,\widehat{\mu}^t) - \mathrm{NSW}(\widehat{p}^t,\mu^*)\big) \\
&\leqslant \sum_{p\in\{p^*,\widehat{p}^t\}} \left|\mathrm{NSW}(p,\mu^*) - \mathrm{NSW}(p,\widehat{\mu}^t)\right| \leqslant 2\alpha^t\sqrt{\frac{2K\log(NKt)}{\gamma \cdot t\epsilon^t}} + \frac{8}{t},
\end{aligned} \tag{8.5}$$

where the penultimate transition holds because $\widehat{p}^t$ is the optimal policy under $\widehat{\mu}^t$, so $\mathrm{NSW}(p^*,\widehat{\mu}^t) \leqslant \mathrm{NSW}(\widehat{p}^t,\widehat{\mu}^t)$, and the final transition follows from Equation (8.4) and the fact that $\mathcal{E}^t \wedge \mathcal{H}^t$ imply $\mathcal{C}_\alpha^t$.

We are now ready to analyze the expected regret of Epsilon-Greedy at round $T$. We have

$$\begin{aligned}
\mathsf{E}[R^T] = \sum_{t=1}^T \mathsf{E}[r^t] &\leqslant \sum_{t=1}^T \mathsf{E}\left[\epsilon^t \cdot 1 + (1-\epsilon^t)\cdot\big(\mathrm{NSW}(p^*,\mu^*) - \mathrm{NSW}(\widehat{p}^t,\mu^*)\big)\right] \\
&\leqslant \sum_{t=1}^T \Bigg(\epsilon^t + \Pr\left[\mathcal{E}^t \wedge \mathcal{H}^t \wedge \mathcal{B}^t\right] \cdot \mathsf{E}\left[\mathrm{NSW}(p^*,\mu^*) - \mathrm{NSW}(\widehat{p}^t,\mu^*) \;\middle|\; \mathcal{E}^t \wedge \mathcal{H}^t \wedge \mathcal{C}_\alpha^t\right] \\
&\qquad\qquad + \Pr\left[\neg\mathcal{E}^t \vee \neg\mathcal{H}^t \vee \neg\mathcal{B}^t\right] \cdot 1\Bigg) \\
&\leqslant \sum_{t=1}^T \left(\epsilon^t + 2\alpha^t\sqrt{\frac{2K\log(NKt)}{\gamma \cdot t\epsilon^t}} + \frac{8}{t} + \frac{4}{t^3} + \frac{1}{t}\right),
\end{aligned}$$

[5] Technically, $n_j^t \geqslant \left\lfloor \frac{b^t}{K} \right\rfloor$ for each arm $j$, but we omit the floor for the ease of presentation.

where the final transition holds due to Equation (8.5), Lemma 24, Lemma 26, and Equation (8.3). Notice that we are using the fact that

$$\Pr[\mathcal{E}^t \wedge \mathcal{H}^t] = \Pr[\mathcal{E}^t] \cdot \Pr[\mathcal{H}^t | \mathcal{E}^t] \geqslant (1 - 2/t^3) \cdot (1 - 2/t^3) \geqslant 1 - 4/t^3.$$

To obtain the first regret bound, we set $\epsilon^t = \Theta\left(N^{\frac{2}{3}} K^{\frac{1}{3}} t^{-\frac{1}{3}} \log^{\frac{1}{3}}(NKt)\right)$ and $\alpha^t = N$, and obtain

$$\mathsf{E}[R^T] = \mathcal{O}\left(N^{\frac{2}{3}} K^{\frac{1}{3}} \log^{\frac{1}{3}}(NKT) \sum_{t=1}^{T} t^{-\frac{1}{3}}\right) = \mathcal{O}\left(N^{\frac{2}{3}} K^{\frac{1}{3}} T^{\frac{2}{3}} \log^{\frac{1}{3}}(NKT)\right).$$

For the second regret bound, we set $\epsilon^t = \Theta\left(N^{\frac{1}{3}} K^{\frac{2}{3}} t^{-\frac{1}{3}} \log^{\frac{2}{3}}(NKt)\right)$ and $\alpha^t = \sqrt{12NK \log(NKt)}$, and obtain

$$\mathsf{E}[R^T] = \mathcal{O}\left(N^{\frac{1}{3}} K^{\frac{2}{3}} \log^{\frac{2}{3}}(NKT) \sum_{t=1}^{T} t^{-1/3}\right) = \mathcal{O}\left(N^{\frac{1}{3}} K^{\frac{2}{3}} T^{\frac{2}{3}} \log^{\frac{2}{3}}(NKT)\right).$$

Note that in both cases, we omit the $\mathcal{O}(1/t)$ and $\mathcal{O}(1/t^3)$ terms because they are dominated by the $\mathcal{O}\left(1/t^{1/3}\right)$ term. In both cases, we use Proposition 9 at the end. □

## 8.6 UCB

**Algorithm 9** UCB

**Input:** Number of agents $N$, number of arms $K$
**Parameters:** Confidence parameter $\alpha^t$ for each $t \in \mathbb{N}$

**for** $t = 1, \ldots, K$ **do**
  $p^t \leftarrow$ policy that puts probability 1 on arm $t$ { Pull arm $t$ deterministically}
**end for**
**for** $t = K + 1, \ldots$ **do**
  Compute the estimated reward matrix $\widehat{\mu}^t$
  $p^t \leftarrow \arg\max_{p \in \Delta^K} \mathrm{NSW}(p, \widehat{\mu}^t) + \alpha^t \sum_{j \in [K]} p_j \cdot r_j^t$, where $r_j^t \triangleq \sqrt{\frac{\log(NKt)}{n_j^t}}$
**end for**

In the classical multi-armed bandit setting, UCB first pulls each arm once. Afterwards, it merges exploration and exploitation cleverly by pulling, in each round, an arm maximizing the sum of its estimated reward and a confidence interval term similar to $r_j^t$ in Algorithm 9. Our multi-agent variant similarly selects a policy that maximizes the estimated Nash social welfare plus a confidence term for a policy, which simply takes a linear combination of the confidence intervals of the arms.

Unlike Explore-First and Epsilon-Greedy, it is not clear if our UCB variant admits an efficient implementation due to the step of computing $\arg\max_{p \in \Delta^K} \mathrm{NSW}(p, \widehat{\mu}) + \alpha^t \sum_{j \in [K]} p_j r_j^t$. Due to the added linear term, the objective is no longer log-concave. This remains a challenging open problem. However, we notice that this can also be viewed as the problem of optimizing a polynomial over a

simplex, which, while NP-hard in general, admits a PTAS when the degree is a constant [de Klerk et al., 2006, 2015]. Hence, when the number of agents $N$ is a constant, this step can be computed approximately, but it remains to be seen how this approximation translates to the regret bound.

We show that UCB achieves the desired $\sqrt{T}$ dependence on the horizon (up to logarithmic factors). In Section 8.3, we show that this is optimal.

**Theorem 47.** *UCB is horizon-independent, and has the following expected regret at any round $T$.*

- *If $\alpha^t = N$ for all $t$, $\mathsf{E}[R^T] = \mathcal{O}\left(NKT^{\frac{1}{2}}\log(NKT)\right)$.*
- *If $\alpha^t = \sqrt{12NK\log(NKt)}$ for all $t$, $\mathsf{E}[R^T] = \mathcal{O}\left(N^{\frac{1}{2}}K^{\frac{3}{2}}T^{\frac{1}{2}}\log^{\frac{3}{2}}(NKT)\right)$.*

*Proof.* Fix one of two parameter choices:

1. $\alpha^t = N$ for all $t$ and $c = N$.
2. $\alpha^t = \sqrt{12NK\log(NKt)}$ for all $t$ and $c = \sqrt{12NK\log(NKT)}$.

Note that in both cases, $\alpha^t \leqslant c$ for all $t$. Hence, $c$ serves as an upper bound on $\alpha^t$ that does not depend on $t$. We show that in both cases, running UCB with the $\alpha^t$ parameter value yields a regret bound of $\mathsf{E}[R^T] = O(cK\sqrt{T}\log(NKT))$. Substituting the two choices of $c$ then yields the two regret bounds.

Let us again focus on the event

$$\mathcal{C}_\alpha^t \triangleq \forall p \in \Delta^K : \left|\mathrm{NSW}(p,\mu^*) - \mathrm{NSW}(p,\widehat{\mu}^t)\right| \leqslant \alpha^t \cdot \textstyle\sum_{j\in[K]} p_j \cdot r_j^t + \dfrac{4}{t}.$$

Recall the clean events $\mathcal{E}^t$ and $\mathcal{H}^t$ defined in Lemmas 24 and 26. Conditioned on $\mathcal{E}^t \wedge \mathcal{H}^t$, note that $\mathcal{C}_\alpha^t$ holds for $\alpha^t = N$ due to Lemma 25, and for $\alpha^t = \sqrt{12NK\log NKt}$ due to Lemma 26. Using Lemmas 24 and 26, and the union bound, we have $\Pr[\neg\mathcal{C}_\alpha^t] \leqslant 1 - \Pr[\mathcal{E}^t \wedge \mathcal{H}^t] = 1 - \Pr[\mathcal{E}^t]\cdot\Pr[\mathcal{H}^t|\mathcal{E}^t] \leqslant 1 - (1-2/t^3)\cdot(1-2/t^3) \leqslant 4/t^3$.

Define a clean event $\mathcal{C}_\alpha^* \triangleq \bigwedge_{t\geqslant\sqrt{T}} \mathcal{C}_\alpha^t$. Here, we do not care about the first $\sqrt{T}$ rounds because the maximum regret from these rounds is $\mathcal{O}\left(\sqrt{T}\right)$, which is permissible given our desired regret bounds. By the union bound, we have $\Pr[\neg\mathcal{C}_\alpha^*] \leqslant T\cdot 4/(\sqrt{T})^3 = 4/\sqrt{T}$. Thus, $\mathcal{C}_\alpha^*$ is a high-probability event. In what follows, we derive an upper bound on the expected regret conditioned on $\mathcal{C}_\alpha^*$, i.e., $\mathsf{E}[R^T|\mathcal{C}_\alpha^*]$. Since conditioning on a high-probability event does not affect the expected value significantly, the desired regret bound will then follow.

For any $t \in [T]$, conditioned on $\mathcal{C}_\alpha^t$ we have that

$$\begin{aligned}\mathrm{NSW}(p^*,\mu^*) \leqslant \mathrm{NSW}(p^*,\widehat{\mu}^t) + \alpha^t \sum_{j\in[K]} p_j^* \cdot r_j^t + \frac{4}{t} &\leqslant \mathrm{NSW}(p^t,\widehat{\mu}^t) + \alpha^t \sum_{j\in[K]} p_j^t \cdot r_j^t + \frac{4}{t} \\ &\leqslant \mathrm{NSW}(p^t,\mu^*) + 2\alpha^t \sum_{j\in[K]} p_j^t \cdot r_j^t + \frac{8}{t},\end{aligned}$$

where the first and the last transition are from conditioning on $\mathcal{C}_\alpha^t$, and the second transition is because $p = p^t$ maximizes the quantity $\mathrm{NSW}(p,\widehat{\mu}^t) + \alpha^t\sum_{j\in[K]} p_j \cdot r_j^t$ in the UCB algorithm.

Let us write $p^{[T]} = (p^1, \ldots, p^T)$ for the random variable denoting the policies used by the algorithm, and $\overline{p}^{[T]} = (\overline{p}^1, \ldots, \overline{p}^T)$ to denote a specific value in $(\Delta^K)^T$ taken by the random variable.

Instead of analyzing $\mathsf{E}[R^T|\mathcal{C}^*_\alpha]$ directly, we further condition on UCB choosing a specific sequence of policies $\overline{p}^{[T]}$. That is, we are interested in deriving an upper bound on $\mathsf{E}[R^T|\mathcal{C}^*_\alpha \wedge p^{[T]} = \overline{p}^{[T]}]$.[6] Interestingly, we show that this quantity is $\mathcal{O}\left(cK\sqrt{T}\log(NKT)\right)$ for *every* possible $\overline{p}^{[T]}$.

Fix an arbitrary $\overline{p}^{[T]}$. For $t \in [T]$ and $j \in [K]$, define $q_j^t = \sum_{s=1}^t \overline{p}_j^s$. Then, $\mathsf{E}[n_j^t|p^{[T]} = \overline{p}^{[T]}] = q_j^t$. For each $j \in [K]$, let $T_j$ be the smallest $t$ for which $q_j^t \geqslant 2\sqrt{T\log(NKT)}$ (if no such $t$ exists, let $T_j = T$); note that given $\overline{p}^{[T]}$, $T_j$ is fixed and not a random variable. Also, we have that $q_j^{T_j} = \Theta\left(\sqrt{T\log(NKT)}\right)$ for each $j \in [K]$.

Let us define a clean event $\mathcal{B} \triangleq \forall j \in [K], n_j^{T_j} \geqslant \sqrt{T\log(NKT)}$. We first show that this is a high probability event. Indeed, using Hoeffding's inequality, we have that for each $j \in [K]$,

$$\begin{aligned}\Pr\left[n_j^{T_j} < \sqrt{T\log(NKT)} \mid \mathcal{C}^*_\alpha \wedge p^{[T]} = \overline{p}^{[T]}\right] &\leqslant \Pr\left[n_j^{T_j} < s_j^{T_j} - \sqrt{T\log(NKT)} \mid \mathcal{C}^*_\alpha \wedge p^{[T]} = \overline{p}^{[T]}\right] \\ &\leqslant \frac{1}{N^2K^2T^2}.\end{aligned}$$

Taking union bound over $j \in [K]$, we have that $\Pr\left[\neg\mathcal{B} \;\middle|\; \mathcal{C}^*_\alpha \wedge p^{[T]} = \overline{p}^{[T]}\right] \leqslant \frac{1}{N^2KT^2}$.

Next, we bound $\mathsf{E}[R^T|\mathcal{C}^*_\alpha \wedge p^{[T]} = \overline{p}^{[T]}]$ by using event $\mathcal{B}$.

$$\begin{aligned}&\mathsf{E}\left[R^T \mid \mathcal{C}^*_\alpha \wedge p^{[T]} = \overline{p}^{[T]}\right] = \sum_{t=1}^T \mathsf{E}\left[\mathrm{NSW}(p^*, \mu^*) - \mathrm{NSW}(p^t, \mu^*) \;\middle|\; \mathcal{C}^*_\alpha \wedge p^{[T]} = \overline{p}^{[T]}\right] \\ &\leqslant \max(K, \sqrt{T}) + \sum_{t=\max(K,\sqrt{T})+1}^T \Bigg(1 \cdot \mathsf{E}\left[\mathrm{NSW}(p^*, \mu^*) - \mathrm{NSW}(p^t, \mu^*) \mid \mathcal{C}^*_\alpha \wedge p^{[T]} = \overline{p}^{[T]} \wedge \mathcal{B}\right] \\ &\qquad + \Pr\left[\neg\mathcal{B} \mid \mathcal{C}^*_\alpha \wedge p^{[T]} = \overline{p}^{[T]}\right] \cdot 1\Bigg) \\ &\leqslant \max(K, \sqrt{T}) + \sum_{t=\max(K,\sqrt{T})+1}^T \mathsf{E}\left[2\alpha^t \sum_{j\in[K]} \overline{p}_j^t \cdot r_j^t + \frac{8}{t} \;\middle|\; \mathcal{C}^*_\alpha \wedge p^{[T]} = \overline{p}^{[T]} \wedge \mathcal{B}\right] \\ &\qquad + T \cdot \Pr\left[\neg\mathcal{B} \mid \mathcal{C}^*_\alpha \wedge p^{[T]} = \overline{p}^{[T]}\right] \\ &\leqslant \max(K, \sqrt{T}) + 1 + 2c\sqrt{2\log(NKT)} \sum_{t=1}^T \sum_{j\in[K]} \frac{\overline{p}_j^t}{\sqrt{c_j}}. \end{aligned} \tag{8.6}$$

The final transition holds because $\alpha^t \leqslant c$ for all $t$, $r_j^t = \sqrt{\frac{2\log(NKT)}{n_j^t}}$, and conditioned on $\mathcal{B}$, $n_j^t \geqslant c_j$ for each $j \in [K]$ and $t \in [T]$, where $c_j = 1$ if $t < T_j$, and $c_j = \sqrt{T\log(NKT)}$ if $t \geqslant T_j$. Hence,

$$\begin{aligned}\mathsf{E}\left[R^T \;\middle|\; \mathcal{C}^*_\alpha \wedge p^{[T]} = \overline{p}^{[T]}\right] &\leqslant \max(K, \sqrt{T}) + 1 + 2c\sqrt{2\log(NKT)} \sum_{j\in[K]} \sum_{t=1}^T \frac{\overline{p}_j^t}{\sqrt{c_j}} \\ &= \max(K, \sqrt{T}) + 1 + 2c\sqrt{2\log(NKT)} \sum_{j\in[K]} \left(\sum_{t=1}^{T_j-1} \frac{\overline{p}_j^t}{1} + \sum_{t=T_j}^T \frac{\overline{p}_j^t}{\sqrt{T\log(NKT)}}\right)\end{aligned}$$

[6] Note that even after conditioning on $p^{[T]} = \overline{p}^{[T]}$, there is still randomness left in sampling actions from the policies and sampling the rewards of those actions.

$$\leqslant \max(K, \sqrt{T}) + 1 + 2c\sqrt{2\log(NKT)} \sum_{j \in [K]} \left( q_j^{T_j} + \frac{T}{\sqrt{T \log(NKT)}} \right) = \mathcal{O}\left( cK\sqrt{T}\log(NKT) \right).$$

Because this bound holds for every possible $\overline{p}^{[T]}$, we also have that $\mathsf{E}[R^T | \mathcal{C}_\alpha^*] = \mathcal{O}\left( cK\sqrt{T}\log(NKT) \right)$. Finally, we can see that

$$\begin{aligned}\mathsf{E}[R^T] &= \Pr[\mathcal{C}_\alpha^*] \cdot \mathsf{E}[R^T | \mathcal{C}_\alpha^*] + \Pr[\neg\mathcal{C}_\alpha^*] \cdot \mathsf{E}[R^T | \neg\mathcal{C}_\alpha^*] \\ &\leqslant 1 \cdot \mathcal{O}\left( cK\sqrt{T}\log(NKT) \right) + \frac{4}{\sqrt{T}} \cdot 1 = \mathcal{O}\left( cK\sqrt{T}\log(NKT) \right).\end{aligned}$$

Recall that substituting $c = N$ and $c = \sqrt{12NK\log(NKT)}$ yields the two regret bounds. □

We emphasize that our analysis of the multi-agent UCB differs significantly from the analysis of the classical (single-agent) UCB. For example, the use of clean event $\mathcal{C}_\alpha^*$ is unique to our analysis. More importantly, the expression in Equation (8.6) is also unique to our setting in which the algorithm can "pull" a probability distribution over the arms. The corresponding expression in case of the classical UCB turns out to be much simpler and straightforward to bound. In contrast, we need to use additional tricks to derive the bound of $\mathcal{O}\left( cK\sqrt{T}\log(NKT) \right)$.

Finally, in the proof presented above, we showed that, assuming the clean event $\mathcal{C}_\alpha^*$, the expected regret is small conditioned on *any* sequence of policies that the UCB algorithm might use. At the first glance, this may seem surprising. However, a keen reader can observe that the clean event $\mathcal{C}_\alpha^*$ can only occur when the UCB algorithm uses a "good" sequence of policies that leads to low expected regret. A similar phenomenon is observed in the analysis of the classical (single-agent) UCB algorithm as well (see, e.g., [Slivkins et al., 2019]): assuming a different clean event, the classical UCB algorithm is guaranteed to not pull suboptimal arms too many times.

## 8.7 Discussion

We introduce a multi-agent variant of the multi-armed bandit problem in which different agents have different preferences over the arms and we seek to learn a tradeoff between the arms that is fair with respect to the agents, where the Nash social welfare is used as the fairness notion. Our work leaves several open questions and directions for future work.

**Computation.** As we observed in the chapter, our Explore-First and Epsilon-Greedy variants can be implemented in polynomial time. However, for our UCB variant, it is not clear if the step of computing $\arg\max_{p \in \Delta^K} \mathrm{NSW}(p, \widehat{\mu}) + \alpha^t \sum_{j \in [K]} p_j r_j^t$ can be computed in polynomial time due to the added linear term at the end. As we mentioned in Section 8.6, there exists a PTAS for this step when the number of agents $N$ is constant, but the complexity in the general case remains open. One might wonder if it helps to take the logarithm of the Nash social welfare, i.e., solve $\arg\max_{p \in \Delta^K} \log \mathrm{NSW}(p, \widehat{\mu}) + \alpha^t \sum_{j \in [K]} p_j r_j^t$. Indeed, since $\log \mathrm{NSW}$ is a concave function, this can be solved efficiently. However, our key lemmas use bounds on the $\mathrm{NSW}$ function that do not hold for $\log \mathrm{NSW}$ function. Further, such an approach would yield a regret bound where the regret is in terms of $\log \mathrm{NSW}$, which cannot be easily converted into regret in terms of $\mathrm{NSW}$. Jones et al. [2023] define a different version of UCB which can be implemented in polynomial time and provides slighly better regret. The main idea is that instead of

applying the confidence term in each policy, this term is applied to each arm and then the policy that achieves the maximum Nash social welfare is chosen.

**Logarithmic regret bound for UCB.** In the classical stochastic multi-armed bandit setting, UCB has two known regret bounds with optimal dependence on $T$. There is an *instance-independent* bound that grows roughly as $\sqrt{T}$ (where the constants depend only on $K$ and not on the unknown mean rewards in the given instance) and an *instance-dependent* bound that grows roughly as $\log T$ (where the constants may depend on the unknown mean rewards in the given instance in addition to $K$). While we recover the former bound in our multi-agent setting, we were not able to derive an instance-dependent logarithmic regret bound. This remains a major challenging open problem.

**Improved lower bounds.** In Section 8.3, we observe that the instance-independent $\Omega(\sqrt{KT})$ lower bound from the classical setting also holds in our multi-agent setting. Given that our upper bounds increase with $N$, it would be interesting to see if we can derive lower bounds that also increase with $N$. Deriving instance-dependent lower bounds in our setting would also be interesting.

**Fairness.** While maximizing the Nash social welfare is often seen as a fairness guarantee of its own, as discussed in the introduction, the policy with the highest Nash social welfare is also known to satisfy other fairness guarantees such as the core [Fain et al., 2018]. However, it is not clear if the additive regret bounds we derive in terms of the Nash social welfare also translate to bounds on the amount by which these other fairness guarantees are violated. Considering other fairness guarantees such as the egalitarian welfare and bounding their total violation is also an interesting direction for the future.

**Multi-agent extensions.** More broadly, our work opens up the possibility of designing multi-agent extensions of other multi-armed bandit problems. For example, one can consider a multi-agent dueling bandit problem, in which an algorithm asks an agent (or all agents) to compare two arms rather than report their reward for a single arm. Meaningfully defining the regret for such frameworks and designing algorithms that bound it is an exciting future direction.

# Chapter 9

# Two-sided Markets Meet Fair Division

## 9.1 Introduction

Popular online platforms such as Facebook, Netflix, and Amazon serve as recommender systems. In a typical recommender system, there is a set of products and a set of customers, and the system recommends to each customer a subset of the products based on their interest. Each customer has a utility function describing the value that they have for being recommended every possible subset of products, and each customer may have a different utility function. Most platforms focus solely on optimizing the experience of the customers, but this can cause products with only slight differences have vastly different rates of exposure. In the extreme hypothetical case, if every customer likes product A ever so slightly more than product B, product A might be recommended to all the customers while product B is recommended to none, despite only slight differences between the two. This can lead to the rich-gets-richer effect of popular products becoming more and more popular while the other products vanishing into oblivion. This reduces competition on the platform, which is harmful for its long-term sustainability. Here, we tackle this problem by requiring that the recommender algorithm be fair not only to the customers, but also to the products at the same time.

Let us formalize this idea a bit more in detail. In a recommender system, each customer can be exposed up to a limited number of products – for example, a web page may have the space to make up to $5$ recommendations to a visiting customer – and producers (who produce the products being recommended) can also have a utility function describing the value that they have for their products being recommended to every possible set of customers. For example, the utility function of a producer may simply measure the number of exposures that their product gets, i.e., the number of customers it is recommended to. If we ignore the customers' preferences, this boils down to a traditional resource allocation problem of allocating goods to agents, where the 'goods' are the customers and the 'agents' are the producers. On the other hand, if we ignore the producers' preferences, then we get again get a traditional resource allocation problem, where now the 'goods' are the products and the 'agents' are the customers.

A central goal in the traditional fair division literature is to find an allocation of the goods to the agents that is *fair*. One desirable notion of fairness is *envy-freeness* [Foley, 1967], which requires that no agent prefer another agent's allocation of goods to her own. This is a compelling definition but, due to the discrete nature of the problem, cannot always be satisfied. Instead, we must consider relaxed

versions, with one popular relaxation being *envy-freeness up to one good* (EF1) [Lipton et al., 2004, Budish, 2011], which requires that any pairwise envy can be eliminated by hypothetically removing a single good from the envied agent's allocation. An allocation satisfying EF1 always exists for a broad class of agent utility functions [Lipton et al., 2004].

As we described above, it can be harmful in the long run to ignore either the customers' (resp., the producers') preferences and achieve fairness with respect to only the producers (resp., the customers). Instead, our goal is to design the recommender system in a way that achieves a fairness notion (such as EF1) with respect to both customers and producers simultaneously, by taking both their preferences into account.

This endeavor of achieving fairness with two-sided preferences is immediately reminiscent of the matching literature. In two sided matching, it is generally assumed that each agent has ordinal preferences over the other side, and a matching is sought that is in some sense *stable* to individual or group deviations. It is well-known that stability is closely related to envy-freeness in the sense that a one-to-one matching is stable if and only if it eliminates *justified envy* [Abdulkadiroğlu and Sönmez, 2003], a requirement specifying that any envy that $i$ may feel for $j$'s match is "justified" by $j$'s match preferring $j$ to $i$. In many-to-one settings, the notions remain tightly connected, with a stable matching being one that eliminates justified envy and has no waste. Justified envy-freeness has also been studied as in its own right in the many-to-one setting [Wu and Roth, 2018, Yokoi, 2020].

Justified envy, and therefore stability, fundamentally rely on the idea that the less an agent is preferred by the other side, the lower her own entitlement should be. However, in some applications, it may be desirable to provide *equal* entitlements or moral claims to agents regardless of how valued they are by the agents on the other side. For example, many customers may prefer the same popular product, but it is not clear that a fair recommendation system should expose this famous product much more than a new product. Therefore, while stability is a valuable notion in many settings, in this work we consider two-sided preferences while incorporating traditional notions from fair division.

**Our contributions.** We introduce and study a two-sided resource allocation setting in which we have two groups of agents, where each agent has preferences over agents on the other side. Each agent must be "matched" to a subset of agents on the other side, subject to a maximum degree constraint. Our goal is to find a many-to-many matching that provides fairness to both sets of agents simultaneously. The standard resource allocation setting is a special case of our model in which each good can be matched to at most one agent, each agent can be matched to any number of goods, and goods are indifferent to which agent they are assigned to.

As a natural tradeoff between expressiveness and succinctness, we restrict our attention to additive preferences, in which the utility for being matched to a group of agents is equal to the sum of utilities for being matched to each agent in the group individually. While conceptually simple, additive preferences have led to a rich body of work in fair division. We focus primarily on the case in which all agents on the same side have the same degree constraint, and the total maximum degree on both sides is equal. In this case, it is reasonable to seek a *complete* matching, which saturates the degree constraints of all the agents on both sides.

We begin by considering *double envy-freeness up to one match* (DEF1), requiring that EF1 hold for both sets of agents simultaneously. We show that a complete matching satisfying (a slight strengthening of) DEF1 does not always exists, but in the special case where both sides have identical ordinal

preferences, it exists and can be computed efficiently using a carefully designed round robin algorithm.

We also ask whether it is possible to find matchings that satisfy *double maximin share guarantee* (DMMS), a two-sided version of the maximin share guarantee. Even when both sides have identical preferences, a complete DMMS matching may not exist, in contrast to the one-sided setting in which an MMS allocation is guaranteed to exist when agents have identical preferences. In general, we show that approximate DMMS and approximate DEF1 are incompatible, although in the special case where the degree constraint is equal to two we can achieve exact versions of both simultaneously.

**Related work.** Most related to our work is that of Patro et al. [2020], who draw on the resource allocation literature to guarantee fairness for both producers and consumers on a two-sided platform. However, in their model, producers are indifferent between the customers; thus, only one side has interesting preferences. Other work [Sühr et al., 2019] has focused on guaranteeing fairness in two-sided platforms over time, rather than in a one-shot setting. Of particular note is the work of Gollapudi et al. [2020], who consider two-sided EF1 in a dynamic setting, but obtain positive results primarily for symmetric binary valuations, a much more restrictive class of valuations than we consider. Tadenuma [2011] studies envy minimization in two-sided matching subject to other notions, including stability, but focuses on ordinal notions of envy and restricts attention to one-to-one matchings.

The theories of matching and fair division each have a rich history. Traditional work in matching theory has focused on one-to-one or many-to-one matchings, beginning with the seminal work of Gale and Shapley [1962] and finding applications in areas such as school choice [Abdulkadiroğlu and Sönmez, 2003, Abdulkadiroğlu et al., 2005, Hatfield et al., 2011], kidney exchange [Saidman et al., 2006], and the famous US resident-to-hospital match.[1] We note that EF1 as a condition becomes vacuous whenever a set of agents has a maximum degree constraint of one, so we focus instead on the more general case of many-to-many matchings. This case has also been well-explored in the matching literature [Roth, 1984, Sotomayor, 1999, Roth and Sotomayor, 1992, Echenique and Oviedo, 2006], although that literature focuses on stability notions, which have a very different flavor to our guarantees.

Our work draws extensively on notions from the fair division literature, particularly envy-freeness and its relaxations [Foley, 1967, Budish, 2011, Lipton et al., 2004] and the maximin share guarantee [Budish, 2011]. Prior work has studied the satisfiability of these properties in the resource allocation setting [Caragiannis et al., 2019b, Procaccia and Wang, 2014, Kurokawa et al., 2016], including the house allocation setting in which each agent is "matched" to a single item [Aigner-Horev and Segal-Halevi, 2022, Beynier et al., 2019, Gan et al., 2019], but, to our knowledge, no work has considered satisfying them on both sides of a market simultaneously.

## 9.2 Preliminaries

For $n \in \mathbb{N}$, define $[n] = \{0, \ldots, n-1\}$. There are two disjoint groups of agents, denoted $N^\ell$ ("left") and $N^r$ ("right"), of sizes $n^\ell$ and $n^r$, respectively. For simplicity of notation, we write $N^\ell = [n^\ell]$ and $N^r = [n^r]$; when referring to an agent by only its index, the group she belongs to will be clear from context. We use indices $i \in [n^\ell]$ and $j \in [n^r]$ to refer to agents on the left and right, respectively. We are given degree constraints $d_i^\ell$ and $d_i^r$ such that each $i \in N^\ell$ and each $j \in N^r$ can be matched to at most $d_i^\ell$

[1] `https://www.nrmp.org/`

and $d_j^r$ agents on the opposite side, respectively. When $d_i^\ell = d_{i'}^\ell$ for any $i, i' \in N^\ell$ (resp. $d_j^r = d_{j'}^r$ for any $j, j' \in N^r$), we denote by $d^\ell$ (resp. $d^r$) the common degree constraint of all agents in $N^\ell$ (resp. $N^r$).

A (many-to-many) *matching* $M$ is represented as a binary $n^\ell \times n^r$ matrix, where $M(i,j) = 1$ if $i \in N^\ell$ and $j \in N^r$ are matched, and $M(i,j) = 0$ otherwise. With slight abuse of notation, we denote $M_i^\ell = \{j \in N^r : M(i,j) = 1\}$ and $M_j^r = \{i \in N^\ell : M(i,j) = 1\}$ as the sets of agents on the opposite side that agents $i \in N^\ell$ and $j \in N^r$ are matched to, respectively. We say that $M$ is *valid* if it respects the degree constraints, i.e., if $|M_i^\ell| \leqslant d_i^\ell$ for each $i \in N^l$ and $|M_j^r| \leqslant d_j^r$ for each $j \in N^r$. Hereinafter, we omit the term valid, but will always refer to valid matchings. We say that $M$ is *complete* if $\sum_{i \in N^\ell} |M_i^\ell| = \sum_{j \in N^r} |M_j^r| = \min(\sum_{i \in N^\ell} d_i^\ell, \sum_{j \in N^r} d_j^r)$. That is, a complete matching is one in which either every agent on the left has their degree constraint met exactly, or every agent on the right does.

Each agent $i \in N^\ell$ has a valuation function $u_i^\ell : N^r \to \mathbb{R}_{\geqslant 0}$ and each agent $j \in N^r$ has a valuation function $u_j^r : N^\ell \to \mathbb{R}_{\geqslant 0}$. When agents $i \in N^\ell$ and $j \in N^r$ are matched, they simultaneously receive utilities $u_i^\ell(j)$ and $u_j^r(i)$, respectively. We assume that utilities are additive. Thus, with slight abuse of notation, the utilities to agents $i \in N^\ell$ and $j \in N^r$ under matching $M$ are $u_i^\ell(M_i^\ell) = \sum_{j \in M_i^\ell} u_i^\ell(j)$ and $u_j^r(M_j^r) = \sum_{i \in M_j^r} u_j^r(i)$, respectively.

Our main constructive results take only the agents' preference orders as input. For agent $i \in N^\ell$ (resp. $j \in N^r$), we denote by $\sigma_i^\ell$ (resp. $\sigma_j^r$) a linear order over $N^r$ (resp. $N^\ell$) which is consistent with the valuation function $u_i^\ell$ (resp. $u_j^r$), i.e., $j \succ_{\sigma_i^\ell} j'$ whenever $u_i^\ell(j) > u_i^\ell(j')$ (resp. $i \succ_{\sigma_j^r} i'$ whenever $u_j^r(i) > u_j^r(i')$).[2] With a slight abuse of notation, we denote with $\sigma_i^\ell(p)$ (resp. $\sigma_j^r(p)$) the position of alternative $p$ in $\sigma_i^\ell$ (resp. $\sigma_j^r$).

Inspired by envy-freeness up to one good (EF1) from classical fair division Budish [2011], Lipton et al. [2004], we define the following fairness guarantee in our setting.

**Definition 18** (Double Envy-Freeness Up To $c$ Matches (DEF$c$))**.** *We say that matching $M$ is envy-free up to $c$ matches (EF$c$) over $N^\ell$ if for each pair of agents $i, i' \in N^\ell$, there exists $S^\ell \subseteq M_{i'}^\ell$ with $|S^\ell| \leqslant c$ such that $u_i^\ell(M_i^\ell) \geqslant u_i^\ell(M_{i'}^\ell \setminus S^\ell)$. Similarly, we say that it is EF$c$ over $N^r$ if, for each pair of agents $j, j' \in N^r$, there exists $S^r \subseteq M_{j'}^r$ with $|S^r| \leqslant c$ such that $u_j^\ell(M_j^r) \geqslant u_j^\ell(M_{j'}^r \setminus S^r)$. We say that $M$ is DEF$c$ if it is EF$c$ over both $N^\ell$ and $N^r$.*

When an algorithm takes as input only the preference rankings, it must ensure that the matching it returns is DEF$c$ for all possible valuation functions which could have induced the rankings. It is easy to observe that this is equivalent to satisfying the following stronger guarantee which uses the stochastic dominance (SD) relation.This is akin to the SD-EF1 strengthening of EF1 [Aziz et al., 2023a].

**Definition 19** (SD Double Envy-Freeness Up To $c$ Matches (SD-DEF$c$))**.** *We say that matching $M$ is SD-envy-free up to $c$ matches (SD-EF$c$) over $N^\ell$ if, for every $t \in [n^r]$,*

$$\textstyle\sum_{p=0}^{t} M(i, \sigma_i^\ell(p)) \geqslant \sum_{p=0}^{t} M(i', \sigma_i^\ell(p)) - c, \forall i, i' \in N^\ell,$$

*and is SD-EF$c$ over $N^r$ if, for every $t \in [n^\ell]$,*

$$\textstyle\sum_{p=0}^{t} M(\sigma_j^r(p), j) \geqslant \sum_{p=0}^{t} M(\sigma_j^r(p), j') - c, \forall j, j' \in N^r.$$

*$M$ is called SD-DEF$c$ if it is SD-EF$c$ over both $N^\ell$ and $N^r$.*

[2] Ties among agents with equal utility are broken arbitrarily.

Finally, we extend a different fairness notion from classical fair division called the maximin share guarantee (MMS).

**Definition 20** ($\alpha$-Double Maximin Share Guarantee ($\alpha$-DMMS)). *Let $\mathcal{M}$ denote the set of valid matchings. The* maximin share value *of agent $i \in N^\ell$ is defined as*

$$MMS_i^\ell = \max_{M \in \mathcal{M}} \min_{i' \in N^\ell} u_i^\ell(M_{i'}^\ell),$$

*and the maximin share value of agent $j \in N^r$ is defined as*

$$MMS_j^r = \max_{M \in \mathcal{M}} \min_{j' \in N^r} u_j^r(M_{j'}^r).$$

*Given $\alpha \in [0, 1]$, matching $M$ is called $\alpha$-maximin share fair ($\alpha$-MMS) over $N^\ell$ if $u_i^\ell(M_i^\ell) \geqslant \alpha \cdot MMS_i^\ell$ for every $i \in N^\ell$, and $\alpha$-MMS over $N^r$ if $u_j^r(M_j^r) \geqslant \alpha \cdot MMS_j^r$ for every $j \in N^r$. It is called $\alpha$-DMMS if it is $\alpha$-MMS for both $N^r$ and $N^r$. When $\alpha = 1$, we write DMMS instead of 1-DMMS.*

The notions of (SD-)DEF1 and DMMS are incomparable to the traditional notions of stability and justified envy-freeness, as the following example shows.

**Example 10.** *Suppose $N^\ell = N^r = \{0, 1, 2, 3\}$, the common degree requirement is 2, and each side has identical ordinal preference $0 \succ 1 \succ 2 \succ 3$ over the other side. The only matching that is stable and eliminates justified envy is the one that matches each $i \in \{0, 1\}$ on the left with every $j \in \{0, 1\}$ on the right, and each $i \in \{2, 3\}$ on the left with every $j \in \{2, 3\}$ on the right. Indeed, if some $i \in \{0, 1\}$ on the left is not matched to some $j \in \{0, 1\}$ on the right, then $j$ must be matched to some $i' \in \{2, 3\}$ on the left, which would make $(i, j)$ a blocking pair, and $i$ would (justifiably) envy $i'$ for her match with $j$. However, this matching violates DEF1 when, for example, agent 2 on the left has more value for agent 1 on the right than for agents 2 and 3 on the right combined (as this would leave her envious of agents 0 and 1 on the left, even after ignoring their match to agent 0 on the right). Note that this matching also violates DMMS, since each agent on the left could partition those on the right into bundles $\{0, 3\}, \{0, 3\}, \{1, 2\}, \{1, 2\}$, guaranteeing themselves a better bundle than the $\{2, 3\}$ that agents 2 and 3 receive. On the other hand, any one-to-one matching satisfies (SD-)DEF1 and DMMS, but many one-to-one matchings are not stable or free of justified envy.*

## 9.3 Double Envy-Freeness Up To One Match

In this section, we focus on double envy-freeness up to one match, more specifically, its strengthening SD-DEF1. We begin in Section 9.3 by presenting an impossibility result that holds even under quite restrictive conditions. Then, we present a positive result for the case that all agents have maximum degree constraint equal to two, and one side has identical preferences. Moreover, we present an algorithm that efficiently computes an SD-DEF1 matching whenever both groups of agents have identical ordinal preferences.

### SD-DEF1 Matchings May Not Exist

Our first main result says that a complete SD-DEF1 matching may not exist. Observe that without the completeness condition an empty matching is trivially SD-DEF1.

**Theorem 48.** *A complete SD-DEF1 matching is not guaranteed to exist.*

*Proof.* Let $d \geqslant 3$ and $n = 4d$. Let all agents in $N^\ell$ have identical preferences over agents in $N^r$ given by $0 \succ \ldots \succ n$. To define the preferences of agents in $N^r$ over $N^\ell$, let us partition the agents in $N^r$ into $d$ blocks: $B^r_m = \{4m, \ldots, 4(m+1)-1\}$ for $m \in [d]$. We define a preference ranking $\rho_m$ for each block $B^r_m$, and let all agents in the block have this preference ranking. The first three rankings $\rho_0$, $\rho_1$, and $\rho_2$ are shown below. The agents not shown in these rankings (marked "remaining agents") can appear in an arbitrary order at the end. Rankings $\rho_3, \ldots, \rho_{d-1}$ can be completely arbitrary.

- $\rho_0 = 0 \succ 1 \succ 2 \succ 3 \succ \ldots$
- $\rho_1 = 0 \succ 1 \succ 4 \succ 5 \succ \ldots$
- $\rho_2 = 2 \succ 3 \succ 4 \succ 5 \succ \ldots$

Once again, we claim that this instance does not admit any complete SD-DEF1 matching. Suppose for contradiction that it does.

Under such a matching, each agent in $N^\ell$ must be matched to exactly one agent in $B^r_m$ for every $m \in [d]$. Suppose for contradiction that this is not true; let $m$ be the smallest index for which the statement fails. Then, because the total degree in $B^r_m$ is $4d = |N^\ell|$, there must exist agents $i, i' \in N^\ell$ such that $i$ is matched to no agent from $B^r_m$, $i'$ is matched to at least two agents from $B^r_m$, and both $i$ and $i'$ are matched to exactly one agent from $B^r_{m'}$ for each $m' < m$. This would violate SD-EF1 with respect to agents $i$ and $i'$.

Notice that each agent in $B^r_0$ must be matched to exactly one agent from $S_0 = \{0, 1, 2, 3\} \subset N^\ell$. If this is not true, then because each agent in $S_0$ is matched to exactly one agent from $B^r_0$ and $|S_0| = |B^r_0| = 4$, we must have agents $j, j' \in B^r_0$ such that agent $j$ is matched to at least two agents from $S_0$ while agent $j'$ is matched to none of them. This would violate SD-EF1 with respect to agents $j$ and $j'$. By a similar reasoning, every agent in $B^r_1$ must be matched to exactly one agent in $S_1 = \{0, 1, 4, 5\} \subset N^\ell$, and every agent in $B^r_2$ must be matched to exactly one agent in $S_2 = \{2, 3, 4, 5\} \subset N^\ell$.

Consider agent $j \in B^r_0$ that agent $4$ from $N^\ell$ is matched to. By the first observation above, there must be a unique agent $i \in S_0$ who is also matched to agent $j$. We take two cases.

**Case 1:** Suppose $i \in \{0, 1\}$. Then, agent $j$ is matched to two agents from $\{0, 1, 4\} \subset S_1$ — agents $i$ and $4$. In contrast, the agent $j' \in B^r_1$ who is matched to agent $5$ from $S_1$ is not matched to any agent from $\{0, 1, 4\} \subset S_1$ by the second claim above. Hence, SD-EF1 is violated with respect to agents $j$ and $j'$, which is a contradiction.

**Case 2:** Suppose $i \in \{2, 3\}$. Then, agent $j$ is matched to two agents from $\{2, 3, 4\} \subset S_2$ — agents $i$ and $4$. In contrast, the agent $j' \in B^r_2$ who is matched to agent $5$ from $S_2$ is not matched to any agent from $\{2, 3, 4\} \subset S_2$ by the third claim above. Hence, SD-EF1 is violated with respect to agents $j$ and $j'$, which is a contradiction. □

The above proof uses a counterexample in which both sides have the same number of agents (that is, $N^\ell = N^r = n$), all agents have the same degree constraint ($d^\ell = d^r = d$), and one group of agents have identical preferences ($u^\ell_i = u^\ell_{i'}$ for all $i, i' \in N^\ell$). Thus, Theorem 48 holds even in this restricted case, and continues to hold for more general settings.

A natural question, that we leave open, is whether the impossibility continues to hold when we relax SD-DEF1 to DEF1. We have found no counterexamples (even for SD-DEF1) via simulation; the counterexample for SD-DEF1 is carefully crafted but relies on the strength of SD-DEF1.

Theorem 48 says that a complete SD-DEF1 matching is not guaranteed to exist even under quite restrictive conditions. It is natural to ask whether there exist any settings in which a complete SD-DEF1 matching can be guaranteed. Below, we consider two different restriction: (a) restricted the degree bound and (b) identical utilities on both sides.

## Restricted Degree Bound

Here, we establish the existence of a complete SD-DEF1 matching by restricting the degree bound (with still only one side having identical preferences).

**Theorem 49.** *When $n^\ell = n^r = n$, $d^\ell = d^r = 2$, and at least one group of agents has identical ordinal preferences, a complete SD-DEF1 matching always exists and can be computed in polynomial time.*

The proof follows from two lemmas; one of which treats the case where $n$ is even, and one where $n$ is odd.

**Algorithm 10** Three-Phase-Round-Robin-I $(n, d, \pi^r)$

*Phase 1:*
**for** $j = 0, \ldots, n/2 - 1$ **do**
  Match agent $j$ on the right to her most preferred agent on the left with no existing matches.
**end for**
*Phase 2:*
**for** $j = n/2, \ldots, n - 1$ **do**
  Match agent $j$ on the right with her most preferred agent on the left which already has one match.
**end for**
*Phase 3:*
**for** $j = n - 1, \ldots, n/2$ **do**
  Let $i'$ be the agent on the left that agent $j$ on the right is already matched to.
  Let $j' \neq j$ be the other agent on the right that $i'$ is matched to ($j'$ must exist).
  Match both $j$ and $j'$ to agent $j$'s most preferred agent on the left who has no existing match.
**end for**

**Lemma 27.** *When $d = 2$ and $n$ is even, Algorithm 10 returns a complete SD-DEF1 matching.*

*Proof.* First, observe that any agent in $B_1^r$, and any agent in $B_2^r$ have one match in $B^\ell$ and one match in $N^\ell \setminus B^\ell$. Denote with $y_j$ and $z_j$ the matches of agent $j \in N^r$ in $B^\ell$ and in $N^\ell \setminus B^\ell$, respectively. Moreover notice that for every agent $j \in B_1^r$ there is an agent $j' \in B_2^r$, such that $y_j = y_{j'}$ and $z_j = z_{j'}$, as in Phase 3 when $j'$ is matched to $z_j$, the same agent is also matched to $j$ if $y_j = y_{j'}$.

We can easily verify that the matching is SD-EF1 over $N^\ell$, as each agent in $N^\ell$ has one match in $B_1^r$, and one match in $B_2^r$. Now, we prove that the matching is also SD-EF1 with respect to the agents in $N^r$. We consider the three following cases.

**Case 1:** $j, j' \in B_1^r$**.** Without loss of generality, we assume that $j < j'$. Then, $y_j \succ_{\sigma_j^r} y_{j'}$, as $j$ chooses before $j'$ in Phase 1. Thus, $j$ cannot envy $j'$ for more than one matches. Moreover, as $j$ has one match in $B^\ell$, and one match $N^\ell \setminus B^\ell$, and as $j'$ prefers $z_{j'}$ to any agent in $N^\ell \setminus B^\ell$ (otherwise she could have chosen an agent in $N^\ell \setminus B^\ell$ as in Phase 1 all of them have no matches), we conclude that $j'$ does envy $j$ for more that one matches.

**Case 2:** $j, j' \in B_2^r$**.** Without loss of generality, we assume that $j < j'$. Then, $y_j \succ_{\sigma_j^r} y_{j'}$, as $j$ appears before $j'$ in Phase 2. Thus, $j$ cannot envy $j'$ for more than one matches. On the other hand, $z_{j'} \succ_{\sigma_{j'}^r} z_j$, as $j'$ appears before $j$ on Phase 3. Hence $j'$ does envy $j$ for more that one matches.

**Case 3:** $j \in B_1^r$ **and** $j' \in B_2^r$**.** If $j$ and $j'$ share the same matches, then obviously they don't envy each other. Otherwise, we denote with $\widehat{j}$ the agent in $B_2^r$ that has the same matches with $j$, and with $\widehat{j'}$ the agent in $B_1^r$ that has the same matches with $j'$. Then, the theorem follows from cases 2 and 3, as the matching is SD-EF1 with respect to $j$ and $\widehat{j}$, and with respect to $j'$ and $\widehat{j'}$. □

**Algorithm 11** Three-Phase-Round-Robin-II $(n, d, \pi^r)$

*Phase 1:*
**for** $j = 0, \ldots, \lceil n/2 \rceil - 1$ **do**
  Match agent $j$ on the right to her most preferred agent on the left with no existing matches.
**end for**
*Phase 2:*
**for** $j = \lceil n/2 \rceil, \ldots, n-2$ **do**
  Match agent $j$ on the right with her most preferred agent on the left which already has one match.
**end for**
Let $i'$ be the agent on the left that agent $\lceil n/2 \rceil - 1$ on the right is matched to.
**if** $i'$ has only one match **then**
  Match agent $n-1$ on the right to agent $i'$ on the left.
**else**
  Match agent $n-1$ on the right to her most preferred agent on the left with no existing matches.
**end if**
Match agent $\lceil n/2 \rceil - 1$ on the right with agent $i'$ on the left with exactly one existing match (such an $i'$ must exist)
*Phase 3:*
**for** $j = n-1, \ldots, \lceil n/2 \rceil$ **do**
  Let $i'$ be the agent on the left that agent $j$ on the right is already matched to.
  Let $j' \neq j$ be the other agent on the right whose only match is to agent $i'$.
  Match both agents $j$ and $j'$ to agent $j$'s most preferred agent on the left who has no existing matches.
**end for**
Match the agent on the right with one existing match, to the agent on the left with one existing match.

**Lemma 28.** *When $d = 2$ and $n$ is odd, Algorithm 11 returns a complete DEF1 allocation.*

*Proof.* Intuitively, the algorithm works in three Phases. Let us divide $N^r$ into two sets: $B_1^r = \{0, ..., \lceil n/2 \rceil - 1\}$, and $B_2^r = \{\lceil n/2 \rceil, ..., n-1\}$. In the first Phase, one-by-one we match agents in $B_1^r$ to their most preferred agent on the left who has no prior matches. We do so in the best-to-worst order over agents in $B_1^r$ according to $\sigma^\ell$. Let $B^\ell$ be the set of agents in $N^\ell$ who now have degree 1; note that $|B^\ell| = \lceil n/2 \rceil$.

In the second Phase, we repeat a similar process, except with agents in $B_2^r \setminus n-1$, still in the best-to-worst order according to $\sigma^\ell$, and by matching them to their most preferred agent in $B^\ell$. Then if $i'$ on the left side, with whom $\lceil n/2 \rceil - 1$ on the right is matched to, has only one match, $n-1$ on the right side is matched with her, otherwise $n-1$ is matched with her best choice in $B^\ell$. At this point, $\lfloor n/2 \rfloor$ agents in $B^\ell$ have degree 2, and only has degree 1, and hence $\lceil n/2 \rceil - 1$ agent on the right is matched to agent in $B^\ell$ with degree 1. So, at the end of this Phase, all agents on the right have degree 1, except for $\lceil n/2 \rceil - 1$ that has degree 2, while agents in $B^\ell$ have degree 2 and agents in $N^\ell \setminus B^\ell$ have degree 0.

In the last Phase, we again consider agents in $B_2^r$, but now in the *worst-to-best* order according to $\sigma^\ell$. Note that each such agent $j$ is already matched to some agent $i' \in B^\ell$, and agent $i'$ is also matched to some agent $j'$ from $B_1^r$. Note that all such $j'$ from $B_1^r \setminus \lceil n/2 \rceil - 1$ has degree 1. Hence, if $j' \neq \lceil n/2 \rceil - 1$ both $j$ and $j'$ to the most preferred agent of $j$ from $N^\ell \setminus B^\ell$, otherwise only $j$ is matched to her. At the end, there is one $j$ in $B_1^r$ , and one agent $i$ in $N^\ell \setminus B^\ell$ that have degree one, and we matched them.

First, observe that any agent in $B_1^r \setminus \lceil n/2 \rceil - 1$, and any agent in $B_2^r$ has one match in $B^\ell$ and one match in $N^\ell \setminus B^\ell$, while $\lceil n/2 \rceil - 1$ has two matches in $B^\ell$. Denote with $y_j$ and $z_j$ the matches of agent $j \in N^r \setminus \lceil n/2 \rceil - 1$ in $B^\ell$ and in $N^\ell \setminus B^\ell$, respectively.

We denote with $j^*$ the agent in $B_2^r$ that share the same match on the left with $\lceil n/2 \rceil - 1$ on the right and with $\widehat{j}^*$ the agent in $B_1^r$ that share the same match on the left with $j^*$. We see that for every agent $j' \in B_2^r \setminus j^*$ there is an agent $j \in B_2^r$, such that $y_j = y_{j'}$ and $z_j = z_{j'}$, as in the third Phase when $j'$ is matched to $z_j$, the same agent is also matched to $j$ if $y_j = y_{j'}$.

We can easily verify that the matching is SD-EF1 over $N^\ell$, as each agent in $N^\ell$ either has one match in $B_1^r$, and one match in $B_2^r$, or has two matches in $B_1^r$ but one of them is with $\lceil n/2 \rceil - 1$.

Now, we prove that the allocation is also SD-EF1 with respect to the agents in $N^r$. We consider the three following cases.

**Case 1:** $j, j' \in B_1^r$**.** This case is similar as Case 1 of the proof of lemma 27, so we omit the details.

**Case 2:** $j, j' \in B_2^r$**.** This case is similar as Case 2 of the proof of lemma 27, so we omit the details.

**Case 3:** $j \in B_1^r$ **and** $j' \in B_2^r$**.** First we know $j$ prefers $y_j$ to any agent in $N^\ell \setminus B^\ell$, and as $j'$ has only one match in $B^\ell$, we conclude that $j$ does not envy $j'$ for more than one matches.

We denote and with $\widehat{j}$ the agent in $B_1^r$ that has the same matches with $j' \in B_2^r \setminus j^*$.

We assume that $j \neq \{\lceil n/2 \rceil - 1, \widehat{j}^*\}$. Then, we know from case 2 that $j'$ does not envy $\widehat{j}$ for more than one match, and similar $j$. If $j = \lceil n/2 \rceil$, then either $j'$ and $j$ share the same matches in $B^\ell$, or $j'$ before line 10 has chosen a better agent rather than the one that $j$ is matched to in this line.

Next, if $j = \widehat{j}^*$ and $j' = j^*$, $j$ and $j'$ share one match, and $j'$ does not envy $j$ for more than one match, while if $j' \neq j^*$ we know that $j'$ prefers $y_{j'}$ to $y_{\widehat{j}^*}$, as $y_{\widehat{j}^*}$ has only one match before line 10. □

*Proof of Theorem 49.* The proof follows from Lemmas 27 and 28. □

## Identical Ordinal Preferences on Both Sides

Here, we consider a different restriction: when agents on both sides have identical preferences, i.e., $\sigma_i^\ell = \sigma_{i'}^\ell$ for all $i, i' \in N^\ell$ and $\sigma_j^r = \sigma_{j'}^r$ for all $j, j' \in N^r$.

**Theorem 50.** *When $n^\ell d^\ell = n^r d^r$ and both groups of agents have identical ordinal preferences, a complete SD-DEF1 matching always exists and can be computed efficiently.*

The proof of Theorem 50 follows from a series of lemmas. We start by focusing on the simple case for which $N^\ell = N^r = n$ and $d^\ell = d^r = d$. Theorem 50 follows from progressively reducing the general case to this simple case.

We denote by $\sigma^\ell$ and $\sigma^r$ the ordinal preferences of the agents in $N^\ell$ and $N^r$, respectively. Without loss of generality, assume that $\sigma^\ell = \sigma^r = 0 \succ \ldots \succ n-1$. We want to find an SD-DEF1 matching under which each agent is matched to exactly $d$ agents on the opposite side. A natural idea is to let agents on one side pick agents on the other side in a round-robin fashion. That is, we construct an ordering $R$ over agents on one side, and these agents take turns according to $R$ in a cyclic fashion with each agent, in her turn, making one match to her most preferred agent (i.e. lowest indexed agent) on the opposite side who has less than $d$ matches so far. A standard argument from classical fair division shows that regardless of the ordering $R$, the resulting matching will be SD-EF1 over over the side that does the picking.[3] However, as the example below shows, not all orderings $R$ lead to a matching that also satisfies SD-EF1 over the other side.

**Example 11.** *Consider the case where $n = 5$ and $d = 2$. Suppose the ordering $R$ has agents on the left choose in the order 0, 1, 2, 3, 4. Then, agent $0$ on the right will be matched to agents $0$ and $1$ on the left, while agent $1$ on the right will be matched to agents $2$ and $3$ on the left. SD-EF1 is violated as agent $1$ significantly envies agent $0$ on the right side.*

We now show that when $R$ is carefully designed, SD-EF1 can also be satisfied over the other side, resulting in SD-DEF1. Algorithm 12 takes as input parameters $a \in [n]$ and $x \in \{d, n-d\}$, and for any choices of these parameters, constructs an ordering $R$ over the agents on (say) the left side. Algorithm 13 then uses this ordering to run the round-robin procedure while respecting the degree constraints. Example 12 demonstrates these algorithms.

**Algorithm 12** Round-Robin-Ordering$(n, a, x)$

1: **for** $i \in [n]$ **do**
2: $\quad R(p) = a + px^{-1} \pmod n$ $\quad$ {$x$ and $n$ coprime, so $x^{-1} \pmod n$ exists}
3: **end for**
4: **return** $R$

**Example 12.** *Consider the same instance as Example 11, with $n = 5$ and $d = 2$. Suppose we choose $a = 3$ and $x = d = 2$. Then the round robin ordering returned by Algorithm 12 is $R(0) = 3 + 0 = 3$ (setting $i = 0$), $R(1) = 3 + 3 = 1$ ($i = 3$), $R(2) = 3 + 1 = 4$ ($i = 1$), $R(3) = 3 + 4 = 2$ ($i = 4$), and*

[3] As observed by Biswas and Barman [2018], the standard round robin algorithm is not EF1 when agents have cardinality constraints, but EF1 is retained provided that agents have identical preferences.

**Algorithm 13** Restricted-Round-Robin-Coprime($n, d$)

1: Choose $a \in \{0, \ldots, n-1\}$ and $x \in \{d, n-d\}$
2: $R =$Round-Robin-Ordering$(n, a, x)$ {Round-robin with ordering $R$ over agents on the left}
3: $M(i, j) = 0, \forall i, j \in [n]$
4: **for** $j \in [n], t \in [d]$ **do**
5: $\quad M(R(j \cdot d + t \pmod n), j) = 1$
6: **end for**
7: **return** $M$

$R(4) = 3 + 2 = 0$ *($i = 2$), with all addition performed mod $n$. That is, agents on the left choose in order* $3, 1, 4, 2, 0$. *This results in the matching* $M_3^\ell = \{0, 2\}$, $M_1^\ell = \{0, 3\}$, $M_4^\ell = \{1, 3\}$, $M_2^\ell = \{1, 4\}$, *and* $M_0^\ell = \{2, 4\}$ *(equivalent to the formula provided directly in Line 5 of Algorithm 13). The fact that this is SD-EF1 over $N^\ell$ is easy to check. Examining the matching, note that* $M_0^r = \{1, 3\} = M_4^\ell$, $M_1^r = \{2, 4\} = M_0^\ell$, $M_2^r = \{0, 3\} = M_4^\ell$, $M_3^r = \{1, 4\} = M_2^\ell$, *and* $M_4^r = \{0, 2\} = M_3^\ell$. *That is, the matchings received by agents on the right are the same as those received by agents on the left, up to a cyclic shift. For this matching, SD-EF1 over $N^\ell$ immediately implies SD-EF1 over $N^r$.*

The next result shows that for any choices of the parameters, the resulting matching is SD-DEF1. The idea of the proof is to show that the structure in Example 12 holds in general: for any allowed choice of $(a, x)$, the set of bundles received by agents on the right is the same as the set of bundles received by agents on the left, thus inheriting SD-EF1 from the fact that the matching is constructed by agents on the left choosing in round robin sequence.

**Lemma 29.** *When $n^\ell = n^r = n$ and $d^\ell = d^r = d$ are coprime and both groups of agents have identical ordinal preferences, Algorithm 13 efficiently computes a complete SD-DEF1 matching.*

*Proof.* To avoid the $\pmod n$ notation in this proof, we will treat integers as belonging to the ring $\mathbb{Z}/n\mathbb{Z}$ of integers modulo $n$. Thus, addition, multiplication, and multiplicative inverses will be modulo $n$. Note that $x^{-1}$ exists because $x \in \{d, n-d\} = \{d, -d\}$ is coprime with $n$.

We claim that the ordering $R$ constructed in Algorithm 12 is a valid ordering over the agents in $N^\ell$. Notice that because $x \in \{d, -d\}$ is coprime with $n$, $(p \cdot x^{-1})_{i \in [n]} = [n]$. Thus, each position in the ordering $R$ is mapped to exactly one agent. Because agents on the left take $d$ turns in a cyclic fashion, it is convenient to think of an *extended* ordering $R$ which is the original $R$ concatenated with itself $d$ times: one can check that this still obeys $R(p) = a + px^{-1}$ for all $p \in [nd]$.

Next, we argue that the matching returned is a valid complete matching. Notice that during the round-robin, $d$ agents on the left that are consecutive in the ordering pick a given agent on the right before moving on to the next lowest indexed agent on the right. Further, each agent on the left gets $d$ turns. Hence, it is easy to see that every agent is matched to exactly $d$ agents on the opposite side.

As mentioned earlier, the fact that the returned matching $M$ is SD-EF1 over $N^\ell$ follows directly from the standard round-robin argument in classical fair division: given any pair of agents $i, i' \in N^\ell$, if we ignored the first turn taken by $i'$, then in each round agent $i$ would get a turn before agent $i'$ does, and hence, would not envy agent $i'$ in the SD sense. It remains to show that $M$ is also SD-EF1 over $N^r$. We show that for each agent $j \in N^r$, there exists an agent $i \in N^\ell$ such that $M_j^r = M_i^\ell$. SD-EF1 over $N^r$ will then follow from SD-EF1 over $N^\ell$ given that $\sigma^\ell = \sigma^r$.

Let us focus on agent $j \in N^r$. Because agents on the right are picked from lowest-indexed to highest-indexed, agent $j$ is picked by the $d$ agents from $N^\ell$ who appear consecutively in the (extended) ordering $R$ at indices $jd + t$ for $t \in [d]$. Given that $R(p) = a + px^{-1}$ for all $p \in [nd]$, we immediately have $M_j^r = \{a + (jd + t)x^{-1} : t = [d]\}$.

Next, let us focus on agent $i \in N^\ell$. If she is matched to some agent $j \in N^r$ in a particular turn, then from the observation above, it must be that $i = a + (jd + t)x^{-1}$ for some $t \in [d]$. Solving this for $j$, we get that $j = ((i - a)x - t)d^{-1}$. Varying $t \in [d]$ in this equation gets us the $d$ agents on the right that agent $i$ is matched to: $M_i^\ell = \{((i - a)x - t)d^{-1} : t = [d]\}$.

To show that for each $j \in N^r$, there exists $i \in N^\ell$ with $M_i^\ell = M_j^r$, we take two cases.

If $x = n - d = -d$, then $x^{-1} = (-d)^{-1} = -d^{-1}$. In this case, it is easy to check that taking $i = j$ suffices as $M_j^r = M_j^\ell = \{a - j - td^{-1} : t = [d]\}$.

If $x = d$, then $M_j^r = \{j + a + td^{-1} : t = [d]\}$, while

$$\begin{aligned} M_i^\ell &= \left\{i - a - td^{-1} : t = [d]\right\} \\ &= \left\{i - a - (d - 1 - t)d^{-1} : t = [d]\right\}. \end{aligned}$$

Notice that $M_j^r$ coincides with $M_{j+2a-d^{-1}+1}^\ell$. □

Algorithm 13 executes round-robin with the left side taking turns, and allows freely choosing $a \in [n]$ and $x \in \{d, n - d\}$ to decide their ordering. Note that if the right side takes turns instead, the algorithm still produces a complete SD-DEF1 matching. However, this extension does not find any new matchings. When $x = n - d$, the matching produced is symmetric ($M_i^\ell = M_i^r$ for all $i \in [n]$), and thus the same regardless of which side takes turns. When $x = d$, the allocations on one side are cyclic shift of the allocations on the other side. Hence, any matching produced by the right side taking turns can also be produced by the left side taking turns with appropriately chosen $(a, x)$.

What about allowing choices of $x$ other than just $d$ and $n - d$? At least for $n = 7$, $d = 3$, and $a = 0$, it is easy to check by hand that no other choices of $x$ produce an SD-DEF1 matching. On the other hand, could it be that some of the $2n$ choices of $(a, x)$ are redundant and lead to the same matching as other choices? The following result shows that in every instance, all $2n$ choices lead to different matchings.

**Proposition 10.** *For any inputs $n$ and $d$ to Algorithm 13, the $2n$ possible choices of $(a, x)$ result in distinct matchings.*

*Proof.* Using the same reasoning as in the proof of Lemma 29, for each $p \in [n]$, we have

$$R(p) = a + p \cdot x^{-1} = \begin{cases} a + p \cdot d^{-1}, & \text{if } x = d, \\ a - p \cdot d^{-1} = a + n - p \cdot d^{-1}, & \text{if } x = n - d. \end{cases}$$

Note that if $(a, x)$ generates an ordering $(a, a_1, \ldots, a_{n-1})$, then $(a, \widehat{x})$ with $\widehat{x} \neq x$ generates the ordering $(a, a_{n-1}, \ldots, a_1)$. Moreover, using different choices of $a$ while fixing the choice of $x$ generates orderings that are cyclically shifted versions of one another.

Let us denote with $R$ and $\widehat{R}$ the orderings returned by Algorithm 12 using $(a, x)$ and $(\widehat{a}, \widehat{x})$, respectively, and let $M$ and $\widehat{M}$ denote the matchings returned by Algorithm 13 under orderings $R$ and $\widehat{R}$, respectively. Suppose for contradiction that $(a, x) \neq (\widehat{a}, \widehat{x})$ but $M = \widehat{M}$.

Let $R = (a, a_1, ..., a_{d-1}, ..., a_{n-1})$. Then, because agents in $N^\ell$ pick agents on the opposite side in a round-robin fashion, we immediately have that $M_0^r = \{a, a_1, ..., a_{d-1}\}$. If $d < n/2$, then $a_{n-1} \notin M_1^r$, while if $d > n/2$, then $a \in M_1^r$. Note that $d = n/2$ is not possible because $d$ and $n$ are coprime. We now consider two cases.

**Case I:** $x = \widehat{x}$**.** As $R$ and $\widehat{R}$ are just cyclically shifted, it is easy to observe that $M_0^r = \widehat{M_0^r}$ implies $a = \widehat{a}$, which is a contradiction.

**Case II:** $x \neq \widehat{x}$**.** Let $R'$ be the ordering obtained with $(a, \widehat{x})$. Then, $R' = (a, a_{n-1}, ..., a_{d-1}, ..., a_1)$. Further, note that $\widehat{R}$ is a cyclically shifted version of $R'$, and we want $\widehat{M_0^r} = M_0^r = \{a, a_1, \ldots, a_{d-1}\}$. It is easy to notice that the only way to obtain this is by setting $\widehat{a} = a_{d-1}$, which will induce $\widehat{R} = (a_{d-1}, ..., a_1, a, a_{n-1}, ..., a_d)$, making $a_{d-1}, \ldots, a_1, a$ the first $d$ agents in the order. If $d < n/2$, then we have $a_{n-1} \in \widehat{M_1^r}$ but $a_{n-1} \notin M_1^r$, while if $d > n/2$, then we have $a \notin \widehat{M_1^r}$ but $a \in M_1^r$. In either case, we have that $M_1^r \neq \widehat{M_1^r}$, which is a contradiction. □

Given Proposition 10, one may be tempted to conjecture that these $2n$ choices generate all complete SD-DEF1 matchings. However, the following example shows that this is not the case, leaving open the question of characterizing the set of all complete SD-DEF1 matchings.

**Example 13.** *Consider the instance with $n = 5$, $d = 2$, and identical ordinal preferences on both sides. Let us focus on the following complete matching $M$.*

- $M_0^\ell = M_0^r = \{0, 2\}$
- $M_1^\ell = M_1^r = \{1, 3\}$
- $M_2^\ell = M_2^r = \{0, 4\}$
- $M_3^\ell = M_3^r = \{1, 4\}$
- $M_4^\ell = M_4^r = \{2, 3\}$

*It is easy to verify that this matching is SD-DEF1. Without loss of generality, we can assume that agents in $N^\ell$ pick agents on the opposite side in a round robin fashion. Then, agents 2 and 3 should be in the last two positions of the round-robin ordering as they are matched to the least preferred agent in $N^r$.*

| | $x = d$ | $x = n - d$ |
|---|---|---|
| $a = 0$ | $[0, 3, 1, 4, 2]$ | $[0, 2, 4, 1, 3]$ |
| $a = 1$ | $[1, 4, 2, 0, 3]$ | $[1, 3, 0, 2, 4]$ |
| $a = 2$ | $[2, 0, 3, 1, 4]$ | $[2, 4, 1, 3, 0]$ |
| $a = 3$ | $[3, 1, 4, 2, 0]$ | $[3, 0, 2, 4, 1]$ |
| $a = 4$ | $[4, 2, 0, 3, 1]$ | $[4, 1, 3, 0, 2]$ |

Table 9.1: Round-robin orderings returned by Algorithm 12 for different choices of $(a, x)$ when $n = 5$ and $d = 2$.

*Table 9.1 shows the round robin orderings produced by all the choices of $(a, x)$. The reader can verify that neither places agents 2 and 3 in the last two positions. Hence, $M$ is not returned by any of the choices.*

The proof of Theorem 50 continues by reducing the case where $n$ and $d$ are not coprime to the coprime case. Letting $g = \gcd(n, d)$, we divide both sides into $g$ sub-groups of $n' = n/g$ agents each. Then, we run Algorithm 13 a total of $g^2$ times to match agents from each sub-group on the left to $d' = d/g$ agents from each sub-group on the right. This matches each agent with exactly $d$ agents from the opposite side. Note that we allow each of the $g^2$ calls to Algorithm 13 to use arbitrary choices of $a$ and $x$. Nonetheless, we show that the resulting complete matching must be SD-DEF1.

**Lemma 30.** *When $n^\ell = n^r = n$, $d^\ell = d^r = d$, and both groups of agents have identical ordinal preferences, a complete SD-DEF1 matching always exists and can be computed efficiently.*

*Proof.* Consider Algorithm 14.

**Algorithm 14** Restricted-Round-Robin($n, d$)

1: $g = \gcd(n, d)$, $n' = n/g$, and $d' = d/g$
2: **for** $k, m \in [g]$ **do**
3: $\quad \widehat{i} = \{n' \cdot k + i : i \in [n']\}$
4: $\quad \widehat{j} = \{n' \cdot m + j : j \in [n']\}$
5: $\quad M(\widehat{i}, \widehat{j})$=Restricted-Round-Robin-Coprime($n', d'$)
6: **end for**
7: **return** $M$

For $k, m \in [g]$, define the sets $B_k^\ell = \{n' \cdot k + i : i \in [n']\}$ and $B_m^r = \{n' \cdot m + j : j \in [n']\}$ as used in Algorithm 14. Note that Algorithm 14 calls Algorithm 13 on each pair $(B_k^\ell, B_m^r)$, and as a result, each agent in $B_k^\ell$ is matched to exactly $d'$ agents in $B_m^r$, and vice-versa.

We want to show that the overall matching $M$ produced by Algorithm 14 is SD-DEF1. Let us first show that it is SD-EF1 over $N^\ell$. Consider two arbitrary agents $i, i' \in N^\ell$. Given that their preference rankings are $0 \succ \ldots \succ n-1$, to show that SD-EF1 holds for these two agents, we need to show that for all $t \in [n]$,

$$\sum_{p=0}^{t} M(i, p) \geqslant \sum_{p=0}^{t} M(i', p) - 1.$$

Fix $t \in [n]$. Because $\{B_m^r : m \in [g]\}$ forms a partition of $N^r = [n]$, there exists a unique $m \in [g]$ such that $t \in B_m^r$. For each $m' < m$, each of agents $i$ and $i'$ are marched to exactly $d'$ agents in $B_{m'}^r$. Hence, $\sum_{p=0}^{n' \cdot m - 1} M(i, p) = \sum_{p=0}^{n' \cdot m - 1} M(i', p) = m \cdot d'$. Thus, it remains to show that

$$\sum_{p \in B_m^r : p \leqslant t} M(i, p) \geqslant \sum_{p \in B_m^r : p \leqslant t} M(i', p) - 1. \tag{9.1}$$

Let $k$ and $k'$ be such that $i \in B_k^\ell$ and $i' \in B_{k'}^\ell$. Let $R$ and $R'$ denote the round-robin orderings constructed in Line 2 of Algorithm 13 when called on $(B_k^\ell, B_m^r)$ and $(B_{k'}^\ell, B_m^r)$, respectively. Let $i''$ be the agent whose position in $R$ matches the position of $i'$ in $R'$. Then, $i''$ is matched to the same set of agents from $B_m^r$ during the call to Algorithm 13 on $(B_k^\ell, B_m^r)$ as $i'$ is matched to during the call to Algorithm 13 on $(B_{k'}^\ell, B_m^r)$. Because Algorithm 13 produces an SD-DEF1 matching, Equation (9.1) holds when $i'$ is replaced by $i''$, and therefore, must also hold for $i'$.

Due to the exact same argument, matching $M$ is also SD-EF1 over $N^\ell$, as desired. □

Now, we are ready to prove Theorem 50.

*Proof of Theorem 50.* Without loss of generality, assume that $n^\ell \leqslant n^r$, and hence, $d^\ell \geqslant d^r$. We add $n^r - n^\ell$ dummy agents with indices $n^\ell, \ldots, n^r - 1$ to the left, so each side has exactly $n^r$ agents. We extend the preferences of the agents on the right so that the dummy agents, indexed higher than the real agents, appear at the bottom of their preference rankings.

Now, we run the Restricted-Round-Robin algorithm (Algorithm 14) with inputs $n^r$ and $d^\ell$. Let $M$ denote the matching returned. Note that while each agent on the left has the intended degree $d^\ell$, each agent on the right has degree $d^\ell$, instead of $d^r$.

Finally, we remove the dummy agents from the left, which reduces the degrees of the agents on the right who were matched to them. Let $\widehat{M} = M([n^\ell], [n^r])$ be the matching $M$ restricted to the real agents. Our goal is to show that $\widehat{M}$ is a complete SD-DEF1 matching for the original problem.

First, note that all agents on the left still have degree $d^\ell$ under $\widehat{M}$. To show that all agents on the right now have degree $d^r$ under $\widehat{M}$, we need to show that they are matched to an equal number ($d^\ell - d^r$) of dummy agents under $M$. Suppose this is not the case. Then, there exist agents $j, j' \in N^r$ such that $j$ is matched to at least two more dummy agents than $j'$ under $M$. It is easy to check that this violates SD-DEF1 of $M$, which is a contradiction because Algorithm 14 returns an SD-DEF1 matching. Thus, after removing the dummy agents, the degree of all agents on the right drop to precisely $d^r$. Hence, $\widehat{M}$ is a complete matching.

To show that $\widehat{M}$ is SD-DEF1, note that it is trivially SD-EF1 over the left side because $M$ is SD-EF1 over this side and the matches of these agents do not change. It is also SD-EF1 over the right side because $M$ is SD-EF1 over this side and exactly $d^\ell - d^r$ least-preferred agents are removed from the allocations of every agent in $N^r$. □

Finally, we turn our attention to the general case in which we drop the constraints $n^\ell = n^r$ and $d_i^\ell = d_j^r = d$. We do however require that $n^\ell \cdot d^\ell = n^r \cdot d^r$. The proof of Theorem 50 uses a trick of adding dummy agents to the side with fewer agents, computing an SD-DEF1 matching as per Lemma 30, and then removing the dummy agents. The key is to show that the removal of dummy agents reduces the degrees of the agents on the opposite side exactly as intended and SD-DEF1 is preserved.

We note that it is possible to extend our constructive result slightly beyond the case of $n^\ell \cdot d^\ell = n^r \cdot d^r$. Without loss of generality, assume that $n^\ell \cdot d^\ell < n^r \cdot d^r$. First, note that in this case, no matching is complete. We can still make the degree of each agent on the left equal to $d^\ell$, but the best we can hope for is that the degrees of agents on the right differ by at most 1, i.e., they are either $\lfloor {}^{n^\ell \cdot d^\ell}/_{n^r} \rfloor$ or $\lceil {}^{n^\ell \cdot d^\ell}/_{n^r} \rceil$.[4] In this case, the trick outlined in Theorem 50 only works when the dummy agents are added to the left side, i.e., if $n^\ell \leqslant n^r$. We conjecture that such an SD-DEF1 matching always exists even when $n^\ell > n^r$, but leave it as an open question.

## 9.4 Double Maximin Share Guarantee

In this section, we focus first on the existence of DMMS matchings, and second on the existence of matchings that are DMMS and SD-DEF1 concurrently.

[4] In case that ${}^{n^\ell \cdot d^\ell}/_{n^r}$ is an integer, we can set this to be $d^r$ and achieve exactly equal degrees on the right side too.

We begin by considering the case where agents on both sides have identical preferences, i.e., $u_i^\ell(j) = u_{i'}^\ell(j)$, for any pair of agents $i, i' \in N^\ell$, and any $j \in N^r$, and similarly $u_j^r(i) = u_{j'}^r(i)$, for any pair of agents $j, j' \in N^r$ and any $i \in N^\ell$. We show the following negative result, which stands in contrast to the one-sided fair division setting in which an MMS allocation is guaranteed to exist when agents have identical preferences.

**Theorem 51.** *A $0.89$-DMMS matching may not exist, even when agents on both sides have identical preferences.*

*Proof.* We denote by $u^\ell$ and $u^r$ the cardinal preferences of the agents in $N^\ell$ and $N^r$ respectively. As the utilities are the same across the agents in the same group, we can define $MMS^\ell = MMS_i^\ell$ for all $i \in [n^\ell]$, and $MMS^r = MMS_j^r$ for all $j \in [n^r]$.

Consider the instance with $n = n^\ell = n^r = 7$ and $d = d^\ell = d^r = 3$, $u^\ell(j) = n - j - 1$ for all $j \in [n]$, and $u^r(i) = n - i - 1$ for all $i \in [n]$. Thus, for any complete matching, $\sum_{i \in N^\ell} u^\ell(M_i^\ell) = \sum_{j \in N^r} u^r(M_j^\ell) = 63$. This means that $MMS^\ell = MMS^r \leqslant 9$, because if all agents receive equal utility then they each get utility $9$. Next, we construct a matching $M$ such that $u^\ell(M_i^\ell) = 9$ for all $i \in [n]$.

Without loss of generality, assume that $M_0^\ell$, $M_1^\ell$, and $M_2^\ell$ all contain agent $0$. Then, we know that agents $1$ and $2$ cannot be contained in these bundles, because then they would have value larger than $9$, implying that some other agent receives utility less than $9$. Without loss of generality, we assume that bundles $M_3^\ell$, $M_4^\ell$, and $M_5^\ell$ contain agent $1$. Now, we observe that $M_6^\ell = \{2, 3, 4\}$, as there is no other way to have $u^\ell(M_6^\ell) = 9$. As $0$ and $2$ can not belong to the same bundle (such a bundle would be valued at least 10), we may assume without loss of generality that agent $2$ is contained in $M_3^\ell$, and $M_4^\ell$. Then, the constraint that $u^\ell(M_3^\ell) = u^\ell(M_4^\ell) = 9$ dictates that $M_3^\ell = M_4^\ell = \{1, 2, 6\}$. With these bundles fixed, it is easy to check that the only $M_5^\ell$ that yields $u^\ell(M_5^\ell) = 9$ is $M_5^\ell = \{1, 3, 5\}$. Lastly, without loss of generality, we may assume that $M_0^\ell = M_1^\ell = \{0, 4, 5\}$, and $M_2^\ell = \{0, 3, 6\}$. Hence, we conclude that the following matching is the only one (subject to permutations of $N^\ell$) that satisfies MMS for agents on the left.

- $M_0^\ell = M_1^\ell = \{0, 4, 5\}$
- $M_2^\ell = \{0, 3, 6\}$
- $M_3^\ell = M_4^\ell = \{1, 2, 6\}$
- $M_5^\ell = \{1, 3, 5\}$
- $M_6^\ell = \{2, 3, 4\}$

Now, consider agents $0 \in N^r$ and $4 \in N^r$. Both are matched to agents $0 \in N^\ell$ and $1 \in N^\ell$, but agent $0 \in N^r$ is matched to agent $2 \in N^\ell$ while agent $4 \in N^r$ is matched to agent $6 \in N^\ell$. Therefore, $u^r(M_0^r) \neq u^r(M_4^r)$ (and this difference persists regardless of permutations of $N^\ell$). It is therefore not the case that every agent on the right receives utility 9; in particular, one agent receives utility 8 or less, producing the approximation ratio $\alpha = 8/9 < 0.89$. □

While a DMMS matching may not exist, even when preferences are identical, we can exploit the algorithms presented in Section 11.4 to obtain an approximation to DMMS.

**Theorem 52.** *When $n^\ell d^\ell = n^r d^r$ and both groups of agents have identical utilities, every complete SD-DEF1 matching $M$ is also $\frac{1}{d^\ell}$-MMS over $N^\ell$, and $\frac{1}{d^r}$-MMS over $N^r$.*

*Proof.* Let $M$ be a complete matching that is SD-EF1 over $N^\ell$. Let $u^\ell$ denote the common utility function for agents in $N^\ell$ which induces (without loss of generality) the preference ranking $0 \succ \ldots \succ n^r - 1$ over $N^r$. We show that $M$ is also $1/d^\ell$-MMS over $N^\ell$. By symmetry, the MMS approximation for $N^r$ also follows.

First, note that each agent in $N^\ell$ must be matched to at least one agent in $\{0, \ldots, \lceil n^\ell/d^r \rceil - 1\} \subset N^r$. This is because if some agent $i \in N^\ell$ is matched to none of these agents, then some other agent $i' \in N^\ell$ must be matched to at least two of them, which would violate SD-EF1 over $N^\ell$.

Let $\widehat{M}$ be a matching that is MMS over $N^\ell$; it is easy to see that under identical cardinal valuations, an MMS matching must exist. Because $(\lceil n^\ell/d^r \rceil - 2) \cdot d^r \leqslant (n^\ell/d^r - 1) \cdot d^r < n^\ell$, there must exist $i \in N^\ell$ such that $\widehat{M}_i^\ell \cap \{0, \ldots, \lceil n^\ell/d^r \rceil - 2\} = \emptyset$. In other words, if $j^* = \arg\max_{j \in \widehat{M}_i^\ell} u^\ell(j)$, then $j^* \geqslant \lceil n^\ell/d^r \rceil - 1$. Notice that $MMS^\ell \leqslant u^\ell(\widehat{M}_i^\ell) \leqslant d^\ell \cdot u^\ell(j^*)$ because $|\widehat{M}_i^\ell| = d^\ell$.

Now, because each agent in $N^\ell$ is matched to at least one agent in $\{0, \ldots, \lceil n^\ell/d^r \rceil - 1\}$, it follows that each agent receives utility at least $u^\ell(j^*) \geqslant 1/d^\ell \cdot MMS^\ell$, as desired. □

We next show an almost-matching upper bound that can be achieved by any SD-DEF1 matching, to complement Theorem 52. In fact, we show a more general result that trades off the approximation to DMMS with the approximation to double envy-freeness.

**Theorem 53.** *There exists an instance with $n^\ell = n^r = n$ and $d^\ell = d^r = d$ in which no matching is simultaneously $\frac{c+2}{d}$-DMMS and SD-DEF$c$ for any $c \in [d]$.*

*Proof.* We present an example in which there is no matching that is concurrently $\frac{c+2}{d}$-MMS and SD-EF$c$ over $N^\ell$.

Let $n$ be divisible by $d$. We consider the instance in which, for all $i \in N^\ell$, $u_i^\ell(j) = d + 1$ for $j \in \{0, ..., n/d - 2\}$, $u_i^\ell(j) = 1$ for $j \in \{n/d - 1, ..., n/d + d - 2\}$, and $u_i^\ell(j) = 0$ for $j \in \{n/d + d - 1, ..., n\}$. Hence, $MMS^\ell$ is equal to $d$, as $n - d$ agents of $N^\ell$ can be matched with the first $n/d - 1$ agents of $N^r$ and take utility equal to $d + 1$, and the remaining $d$ agents can have $d$ matches in the set $\{n/d - 1, ..., n/d + d - 2\} \subset N^r$ and receive utility equal to $d$. Because there are less than $n + d^2$ matches available in the subset $\{0, ..., n/d + d - 2\} \subset N^r$, and $n \geqslant d^2$, it must be the case that there are agents on the left that have at most one match to this subset of $N^r$. So, to achieve a SD-EF$c$ allocation with respect to the agents in $N^\ell$, each agent can have at most $c + 1$ matches to this subset. But then only $n - d$ out of $n$ agents have utility at least $d + 1$, and the remaining ones have utility at most $c + 1$. So, there are agents that receive utility at most $c + 1$ while $MMS^\ell = d$. □

Finally, we show that a strong impossibility persists even if we only require SD-EF1 on one side and MMS on the other.

**Theorem 54.** *A matching that satisfies SD-EF1 over $N^\ell$ and MMS over $N^r$ is not guaranteed to exist, even when agents on both sides have identical preferences.*

*Proof.* We consider the instance in which $n^\ell = n^r = n = 11$, and $d^\ell = d^r = d = 3$, while $u^\ell = u_i^\ell(j) = n - j - 1$ for all $i \in N^\ell$ and $j \in N^r$ and $u^r = u_j^r(i) = n - i - 1$ for all $i \in N^\ell$ and

$j \in N^r$. Finding $MMS^\ell$ can be done using the following integer linear program $ILP_1(n, d, v^\ell)$:

$$\begin{array}{lll} \text{maximize} & MMS^\ell & \\ \text{subject to} & \sum_{j\in[n]} M(i,j) \cdot v^\ell(j) \geqslant MMS^\ell, & \forall i \in [n] \\ & \sum_{j\in[n]} M(i,j) = d, & \forall i \in [n] \\ & \sum_{i\in[n]} M(i,j) = d, & \forall j \in [n] \\ & M(i,j) \in \{0,1\}, & \forall i,j \in [n] \end{array}$$

The first constraint ensures that the matchings is MMS over $N^\ell$, and the remaining ones ensure that the matching is valid and complete.

Detecting whether it exists a matching that is SD-EF1 over $N^r$ and MMS over $N^\ell$ can be done using the following integer linear program $ILP_2(n, d, v^\ell, MMS^\ell)$:

$$\begin{array}{ll} \sum_{i\in[r]} M(i,j) \geqslant \sum_{i\in[r]} M(i,j') - 1, & \forall j, j' \in [n], \forall r \in [n] \\ \sum_{j\in[n]} M(i,j) \cdot v^\ell(j) \geqslant MMS^\ell, & \forall i \in [n] \\ \sum_{j\in[n]} M(i,j) = d, & \forall i \in [n] \\ \sum_{i\in[n]} M(i,j) = d, & \forall j \in [n] \\ M(i,j) \in \{0,1\}, & \forall i,j \in [n] \end{array}$$

The first constraint ensures that the matchings is SD-EF1 over $N^r$, the second one that it is MMS over $N^\ell$, and the remaining ones ensure that the matching is valid and complete.

Using the above ILP programs, we check that the aforementioned instance does not admit a matching that is concurrently SD-EF1 over $N^r$ and MMS over $N^\ell$.

□

## 9.5 Simulations

Above, we show that Algorithm 13 finds an SD-DEF1 matching when *both sides* have identical ordinal preferences. Here, we conduct experiments to see how well the algorithm works beyond this restricted case. Specifically, we simulate instances in which we start from identical preferences on both sides but, for some agents in $N^\ell$ and/or some agents in $N^r$, replace their preferences by those drawn from the uniform distribution (impartial culture). We measure the number of agents for whom the resulting matching is not EF1 and use the classic round robin (CRR) algorithm as benchmark.

First, we observe that when even a few agents in $N^\ell$ — the side that picks — have heterogeneous preferences, both Algorithm 13 and CRR have almost similar performance. This is quite intuitive as the proof of Lemma 29 relies on a delicate mathematical formula for the set of agents picked by each agent in $N^\ell$, and even a single agent in $N^\ell$ not picking as expected breaks this.

Next, we restrict our attention to instances where the agents in $N^\ell$ and $p\%$ of agents in $N^r$ have identical preferences. In this case, Algorithm 13 is guaranteed to satisfy SD-EF1 for all agents except

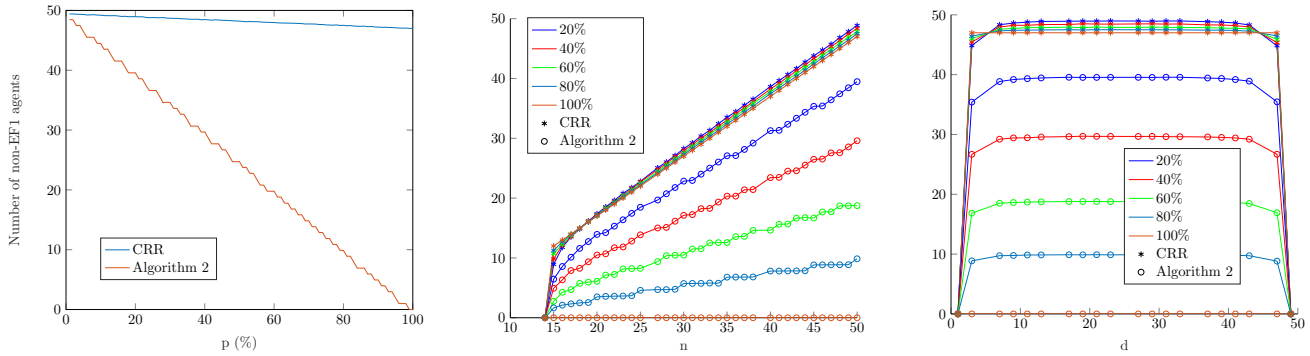


Figure 9.1: Non-EF1 agents for (a) $n = 50$, $d = 23$ and various values of $p$ (left), (b) $d = 13$ and various values of $n$ and $p$ (middle), and (c) $n = 50$ various values of $d$ and $p$ (right).

the $(100 - p)\%$ agents in $N^r$ with their preferences drawn from the uniform distribution. Figure 9.1 shows the number of agents for whom EF1 is violated for different values of $p$ when (a) $n = 50$ and $d = 27$ (left), (b) $n \in \{14, \ldots, 50\}$ and $d = 13$ (middle), and (c) $d \in \{1, \ldots, 49\}$ and $n = 50$ (right).[5] Each point is an aggregation over 1000 random instances.

First, we observe that when $p$ is close to $0\%$, CRR and Algorithm 13 have similar performance, but as $p$ increases, the performance of Algorithm 13 improves substantially (achieving SD-DEF1 when $p = 100\%$) while that of CRR improves negligibly. Quite interestingly, the number of non-EF1 agents does not change significantly with $d$ (except for extremely small or large values of $d$).

We also conducted simulations to compare the two-sided fair division setting in which both sides have preferences and we seek EF/approximate MMS for both sides to the one-sided fair division setting in which only one side has preferences and we want EF/approximate MMS only for that side. We vary $n \in \{3, \ldots, 8\}$ and $d \in \{1, \ldots, n-1\}$, and, for each combination, generate 500 random instances with additive utilities where the utility of each agent for each good/agent is selected uniformly at random from $\{0, \ldots, 20\}$.

For envy, we find a sharp contrast between one-sided EF and two-sided DEF for $n \geqslant 6$ and $2 \leqslant d \leqslant n - 2$: while EF is almost always achievable (in at least $92.6\%$ of instances) while DEF is almost always unachievable (achievable in at most $7.6\%$ instances).

However, for maximin share guarantee, we do not find such a contrast. One-sided (exact) MMS was always achievable in all our simulations. Two-sided (exact) DMMS was also achievable in at least $92.6\%$ instances, and in the remaining, 0.99-DMMS was still achievable.

## 9.6 Discussion

We have introduced a model that bridges two-sided matching and fair division by requiring fairness on both sides of a matching market. We have shown that SD-EF1 can be achieved for agents on both sides, when all agents on the left side and all agents on the right share a common ordinal preference ranking over agents on the other side. When this condition is not satisfied, there may exist no matching that satisfies SD-DEF1. We have also shown that there may not exist a doubly MMS matching even when

[5] We restrict attention to coprime $(n, d)$, so some entries in the plots are skipped.

agents have identical preferences. While we do not rule out a good approximation to DMMS, we show that it is essentially impossible to obtain a good approximation to DMMS if one also requires SD-DEF1.

It is interesting to note that the proofs of Theorems 52 and 53 do not rely on the constraints that an agent in $N^\ell$ can have up to $d$ matches, and can be matched with an agent in $N^r$ at most once. Therefore, these theorems also hold in a one-sided fair division problem where there are $n$ agents and $n/d$ items with $d$ copies each, and all the agents have identical preferences.

Many interesting avenues for future research remain. For example, one can hope to derive weaker positive results in the case where one side has identical preferences. When the side with identical preferences does the picking, Algorithm 13 remains EF1 for that side. In Section 9.5, we conduct empirical simulations and observe that Algorithm 13 remains EF1 for some of the agents on the other side as well (and does better than the classical round-robin algorithm in this aspect). It would also be interesting to compare the two-sided fair division setting with its one-sided counterpart (where only one side has preferences and we seek fairness only for this side). In our simulations, we observe that there is a sharp contrast for envy-freeness (one-sided EF is almost always achievable while two-sided DEF almost always isn't). For the maximin share guarantee, however, there is no contrast: both one-sided MMS and two-sided DMMS are almost always achievable. Theoretically analyzing the probability of satisfiability of these notions in random instances would be an interesting direction for the future. One can also consider two-sided versions of other fairness notions, including those that remain interesting when the degree constraint is 1, which could yield further interesting results in the one-to-one or many-to-one settings. Finally, it would also be interesting to derive positive results when each agent can have a different degree constraint.

# Chapter 10

# Group Fairness in Peer Review

## 10.1 Introduction

Computer Science is a rapidly advanced field, and therefore peer-reviewed conferences are at the heart of the research progress, since their reviewing time is usually quite fast [Halpern and Parkes, 2011, Vrettas and Sanderson, 2015]. In many of these conferences, such as AAMAS, AAAI, and NeurIPS, the assignment of the papers to reviewers is usually an automated procedure, due to their massive scale. Famous automated systems that are used in practise are the Toronto Paper Matching System [Charlin and Zemel, 2013], Microsoft CMT[1], and OpenReview[2]. The authors of the submissions are usually very interested to receive useful feedback from their peers, regarding how they could improve their paper [Travis and Collins, 1991, Nickerson, 2005, Shah, 2022]. Naturally, the overall experience of an author for a peer review procedure highly depends on the quality of the reviews that her manuscripts receive.

In many large conferences, the typical procedure of selecting the reviewers of each manuscript is the following one. First, for each paper-reviewer pair is calculated a similarity score based on various parameters such as the subject area of the paper and the reviewer, the bidding of the reviewer, etc. [Charlin and Zemel, 2013, Mimno and McCallum, 2007, Liu et al., 2014, Rodriguez and Bollen, 2008, Tran et al., 2017]. Then, an assignment is calculated through an optimization problem where the usual objectives are either to maximize the utilitarian social welfare, which is equal to the total similarity, or the egalitarian social welfare, which is equal to the minimum score of each submission, subject to constraints related to the total number of papers that each reviewer can review and the total number of reviewers that each paper should be assigned to. Under both objectives, it is possible that the review quality on some papers is sacrificed. To see that, consider the case that there are four submissions, $p_1$, $p_2$, $p_3$ and $p_4$, and four reviewers, $1$, $2$, $3$ and $4$ who can review up one paper each.

Assume that the first two reviewers have similarity score equal to $1$ for $p_1$ and $p_2$ and equal to $0.9$ for $p_3$ and $p_4$, while the other two reviewers have similarity score equal to a negligible quantity $\epsilon > 0$ for $p_1$ and $p_2$ and have zero similarity score for $p_3$ and $p_4$. This may happen, when $1$ and $2$ work on topics that these papers consider, while $3$ and $4$ belong in a different community. If the goal is to maximize the utilitarian social welfare, then $p_1$ and $p_2$ are assigned to the first two reviewers, while $p_3$'s and $p_4$'s

[1] `https://cmt3.research.microsoft.com/`

[2] `https://github.com/openreview/openreview-matcher`

utilities are completely sacrificed, while if the goal is to maximize the egalitarian social welfare, the opposite happens. Papers that are assigned to inappropriate reviewers may receive poor feedback or even may be unfairly rejected, which may cause their authors to be significantly unsatisfied with the whole procedure. Thus, finding reviewing assignments that are fair is very important, and the last years researchers have focused in this direction [Shah, 2019].

Peng et al. [2022] recently mentioned that a major problem with the prestigious mega conferences is that they constitute the main venues for several communities, and as a result, in some cases, people are asked to review submissions that are beyond their main areas of work. They claim that a reasonable solution is to move to a de-centralized publication process by creating more specialized conferences appropriate for different communities. In particular, they say that by this way "*Reviewers and reviewees will be peers, collaborators, and problem-specific interlocutors, not generic members of a large anonymized community.*". However, this solution could cause the isolation of different communities which in its turn could cause various other problems such as the difficulty of emerging interdisciplinary ideas. Moreover, it usually takes several years until a conference becomes famous and acceptable across the members of a community. So, a reasonable question is

> ...how can we treat each group of researchers in a fair way in the current review and publication processes?

To answer this question, we use the concept of fairness, which, to the best of our knowledge, we are the first that introduce in a peer review setting, called *core* [Gillies, 1953]. In this context, this notion requires that given an assignment there is no subset of authors– who can also serve as reviewers– that can deviate as following: They can find an assignment of their submissions among themselves such that

- no author reviews her own submissions,
- each paper is reviewed by as many reviewers as in the given assignment,
- each reviewer reviews no more papers than in the given assignment, and
- the submissions of each author are assigned to better reviewers than in the given assignment.

Intuitively, this notion of fairness requires that any group of authors is treated in a way that it does not have any incentive to deviate from the given assignment and create its own assignment that meets the constraints of the peer review procedure. In other words, any sub-community in a big conference is treated in a way that it does not have any incentive to deviate from the conference and create its own smaller conference. Note that this definition provides fairness to *every* sub-community and not only to pre-defined ones, and as result it guarantees that even emerged interdisciplinary subcommunities, are treated fairly.

## Our Contribution

In this work, we consider the case that each submission is authored by one agent that also serves as reviewer. A reviewing assignment is valid if each paper is reviewed by $k_p$ reviewers, each reviewer reviews up to $k_a$ papers and no agent reviews her own submissions. To ensure that a valid assignment always exists, we assume that the maximum number of papers that each agent can submit is at most $\lfloor k_a/k_p \rfloor$.

In Section 10.3, we present an efficient algorithm that always returns a valid assignment in the core under very minor assumptions regarding the preferences of the authors for different potential reviewers. In particular, we assume that each author holds an ordinal preference over the reviewers with respect to each of her submissions and the extension of these preferences to preferences over sets of reviewers that review her submissions follows some very natural properties.

In Section 10.4, we show that there are instances where no assignment in the core can provide an approximation better than $\Omega(n)$ with respect to the utilitarian social welfare and bounded approximation with respect to the egalitarian social welfare. Moreover, we show that it is NP-hard to find an assignment in the core with maximum utilitarian social welfare and an assignment in the core that provides bounded approximation to the best egalitarian welfare that can be achieved by any assignment in the core.

In Section 10.5, we conduct experiments with real data and observe that our algorithm achieves good utilitarian and egalitarian social welfare in the average case, while broadly applied methods fail to find assignments in the core, and as a result communities are incentivized to deviate.

## Related Work

The reviewing assignment problem has been extensively studied [Wang et al., 2008]. Toronto Paper Matching System [Charlin and Zemel, 2013] which is a very broadly applied method focuses on maximizing the utilitarian welfare and this approach has been adopted by other popular conference management systems such as EasyChair[3] and HotCRP [4] [Stelmakh et al., 2019]. O'Dell et al. [2005] got a different approach where the goal is maximize the minimum total utility that a paper gets, and Stelmakh et al. [2019] generalized this approach by maximizing the minimum paper score, then maximizing the next smallest paper score, etc. One of the key issues in reviewer assignment is to ensure that the assignment is fair and efficient for the reviewers as well as the papers/authors. Several papers have examined this issue in different respects (see, e.g., [Payan and Zick, 2022, Garg et al., 2010, Lian et al., 2018, Kobren et al., 2019, Shah, 2022]). The core property we focus on can also be viewed as a fairness or efficiency requirement.

Assignment of papers to reviewers is essentially a matching problem and hence has connections with several classical problems in matching. Our model is related to exchange problems with endowments. Agents can be viewed as being endowed by their own papers which they wish to exchange with other agents. In contrast to classical exchange problem with endowments, our model has a distinctive requirement that agents need to give away *all* their items/papers as the papers need to be reviewed by the agent who gets the paper. The difference is crucial as explained next.

A basic exchange problem the Shapley-Scarf housing market in which each agent owns one house. Shapley and Scarf [1974] showed that a simple yet elegant mechanism called *Gale's Top Trading Cycle (TTC)* finds an allocation which is in the core. TTC is based on multi-way exchanges of houses between agents. Since the basic assumption in the model is that agents have strict preferences over houses, TTC is also strict core selecting and therefore Pareto optimal. Our model involves agents getting multiple items. For problems with multiple endowments, Konishi et al. [2001] showed that the core can be empty under additive valuations. Note our problem is different as individual rationality has no bite in our context and an agent is required to give away all of his 'resources' (own papers).

[3] https://easychair.org

[4] https://hotcrp.com

Our model also has connections with matching with two-sided preferences where agents have preferences over reviewers. In many to many matchings with two-sided preferences, several solution concepts have been used to identify desirable matchings. The classical concept of *pairwise stability* requires that there are no two agents who are not partners, but by becoming partners, possibly dissolving some of their partnerships to remain within their quotas and possibly keeping other ones, can both obtain a strictly preferred set of partners. A matching is *corewise-stable* if there is no subset of agents who by forming all their partnerships only among themselves, can all obtain a strictly preferred set of partners. A matching $\mu$ will be called *setwise-stable (SW)* if there is no subset of agents who by forming new partnerships only among themselves, possibly dissolving some partnerships of $\mu$ to remain within their quotas and possibly keeping other ones, can all obtain a strictly preferred set of partners [Sotomayor, 1999]. For a detailed taxonomy of stability concepts for matching for many to many matchings, see the paper by Klaus and Walzl [2009].

## 10.2 Model

For $q \in \mathbb{N}$, define $[q] \triangleq \{1, \ldots, q\}$. There is a set of agents $N = [n]$. Each agent $i$ submits a set of papers $P_i = \{p_{i,1}, \ldots, p_{i,m_i}\}$ for review by her peers, where $m_i \in \mathbb{N}$, and is available to review the submissions of her peers. We refer to $p_{i,\ell}$ as the $\ell$-th submission of agent $i$; when considering the special case of each agent $i$ having a single submission, we will drop $\ell$ and simply write $p_i$. Let $P = \cup_{i \in N} P_i$ be the set of all submissions and $m = \sum_{i \in N} m_i$ be the total number of submissions.

**Assignment.** Our goal is to produce a *(reviewing) assignment* $R : N \times P \to \{0, 1\}$, where $R(i, j) = 1$ if agent $i \in N$ is assigned to review submission $j \in P$. With slight abuse of notation, let $R_i^a = \{j \in P : R(i, j) = 1\}$ be the set of submissions assigned to agent $i$ and $R_j^p = \{i \in N : R(i, j) = 1\}$ be the set of agents assigned to review submission $j$. We want the assignment to be *valid*, i.e., satisfy the following constraints:

- Each agent must be assigned at most $k_a$ submissions for review, i.e., $|R_i^a| \leqslant k_a, \forall i \in N$.
- Each submission must be assigned to $k_p$ agents, i.e., $|R_j^p| = k_p, \forall j \in P$.
- No agent should review one of her own submissions, i.e., $R(i, p_{i,\ell}) = 0, \forall i \in N, \ell \in [m_i]$.

To ensure that a valid assignment always exists, we impose the constraint that $m_i \cdot k_p \leqslant k_a$ for each $i \in N$, which implies that $m \cdot k_p \leqslant n \cdot k_a$. Intuitively, this demands that each agent submitting papers be willing to provide as many reviews as the number of reviews assigned to the submissions of any single agent. For further discussion on this condition, see Section 10.6.

Note that given $N' \subseteq N$ and $P_i' \subseteq P_i$ for each $i \in N'$ with $P' = \cup_{i \in N'} P_i'$, the validity requirements above can also be extended to a restricted assignment $\widehat{R} : N' \times P' \to \{0, 1\}$. Hereinafter, we will assume validity unless specified otherwise or during the process of building an assignment.

**Preferences.** Each agent $i \in N$ has a preference ranking, denoted $\sigma_{i,\ell}$, over the agents in $N \setminus \{i\}$ for reviewing her $\ell$-th submission $p_{i,\ell}$.[5] These preferences can be based on a mixture of many factors, such as how qualified the other agents are to review submission $p_{i,\ell}$, how likely they are to provide

[5]Our algorithms continue to work with weak orders; one can arbitrarily break ties to convert them into strict orders before feeding them to our algorithms.

a positive review for it, etc. Let $\sigma_{i,\ell}(i')$ be the position of agent $i' \in N \setminus \{i\}$ in the ranking. We say that agent $i$ prefers agent $i'$ to agent $i''$ as a reviewer for $p_{i,\ell}$ if $\sigma_{i,\ell}(i') < \sigma_{i,\ell}(i'')$. Again, in the special case where the agents have a single submission each, we drop $\ell$ and just write $\sigma_i$. Let $\vec{\sigma} = (\sigma_{1,1}, \ldots, \sigma_{1,m_1}, \ldots, \sigma_{n,1}, \ldots, \sigma_{n,m_n})$.

While our algorithm takes $\vec{\sigma}$ as input, to reason about its guarantees, we need to define agent preferences over assignments by extending $\vec{\sigma}$. In particular, an agent is assigned a set of reviewers for each of her submissions, so we need to define her preferences over sets of sets of reviewers. First, we extend to preferences over sets of reviewers for a given submission, and then aggregate preferences across different submissions. Instead of assuming a specific parametric extension (e.g., additive preferences), we allow all possible extensions that satisfy two mild constraints; the group fairness guarantee of our algorithm holds with respect to any such extension.

*Extension to a set of reviewers for one submission:* Let $S \succ_{i,\ell} S'$ (resp., $S \succeq_{i,\ell} S'$) denote that agent $i$ strictly (resp., weakly) prefers the set of agents $S$ to the set of agents $S'$ for her $\ell$-th submission $p_{i,\ell}$. We require only that these preferences satisfy the following mild axiom.

**Definition 21** (Order Separability)**.** *For every disjoint $S_1, S_2, S_3 \subseteq N$ with $|S_1| = |S_2| > 0$, if it holds that $\sigma_{i,\ell}(i') < \sigma_{i,\ell}(i'')$ for each $i' \in S_1$ and $i'' \in S_2$, then we must have $S_1 \cup S_3 \succ_{i,\ell} S_2 \cup S_3$.*

An equivalent reformulation is that between any two sets of reviewers $S$ and $T$ with $|S| = |T|$, ignoring the common reviewers in $S \cap T$, if the agent strictly prefers every (even the worst) reviewer in $S \setminus T$ to every (even the best) reviewer in $T \setminus S$, then the agent must strictly prefer $S$ to $T$.

**Example 14.** *Consider the common example of additive preferences, where each agent $i$ has a utility function $u_{i,\ell} : N \setminus \{i\} \to \mathbb{R}_{\geqslant 0}$ over individual reviewers for her $\ell$-th submission, inducing her preference ranking $\sigma_{i,\ell}$. In practice, these utilities are sometimes called similarity scores. Her preferences over sets of reviewers are defined via the additive utility function $u_{i,\ell}(S) \triangleq \sum_{i' \in S} u_{i,\ell}(i')$. It is easy to check that for any disjoint $S_1, S_2, S_3$ with $|S_1| = |S_2| > 0$, $u_{i,\ell}(i') > u_{i,\ell}(i'')$ for all $i' \in S_1$ and $i'' \in S_2$ would indeed imply $u_{i,\ell}(S_1 \cup S_3) > u_{i,\ell}(S_2 \cup S_3)$. Additive preferences are just one example from a broad class of extensions satisfying order separability.*

*Extension to assignments.* Let us now consider agent preferences over sets of sets of reviewers, or equivalently, over assignments. Let $R \succ_i \widehat{R}$ (resp., $R \succeq_i \widehat{R}$) denote that agent $i$ strictly (resp., weakly) prefers assignment $R$ to assignment $\widehat{R}$. Note that these preferences collate the submission-wise preferences $\succ_{i,\ell}$ across all submissions of the agent. We require only that the preference extension satisfies the following natural property.

**Definition 22** (Consistency)**.** *For any assignment $R$, restricted assignment $\widehat{R}$ over any $N' \subseteq N$ and $P' = \cup_{i \in N'} P'_i$ (where $P'_i \subseteq P_i$ for each $i \in N'$), and agent $i^* \in N'$, if it holds that $R^p_{p_{i^*,\ell}} \succeq_{i^*,\ell} \widehat{R}^p_{p_{i^*,\ell}}$ for each $p_{i^*,\ell} \in P'_i$, then we must have $R \succeq_i \widehat{R}$.*

In words, if an agent weakly prefers $R$ to $\widehat{R}$ for the set of reviewers assigned to each of her submissions individually, then she must prefer $R$ to $\widehat{R}$ overall.

**Example 15.** *Let us continue with the previous example of additive utility functions. The preferences of agent $i$ can be extended additively to assignments using the utility function $u_i(R) = \sum_{p_{i,\ell} \in P} u_{i,\ell}(R^p_{p_{i,\ell}})$. It is again easy to check that if $u_{i,\ell}(R^p_{p_{i,\ell}}) \geqslant u_{i,\ell}(\widehat{R}^p_{p_{i,\ell}})$ for each $p_{i,\ell}$, then $u_i(R) \geqslant u_i(\widehat{R})$. Hence,*

**Algorithm 15** CoBRA

**Input:** $N, P, \vec{\sigma}, k_a, k_p$
**Output:** $R$

$R, L, U =$PRA-TTC$(N, P, \vec{\sigma}, k_a, k_p)$
**if** $|U| > 0$ **then**
  $R =$Filling-Gaps$(N, P, \vec{\sigma}, k_a, k_p, R, L, U)$
**end if**

**Algorithm 16** PRA-TTC

**Input:** $N, P, \vec{\sigma}, k_a, k_p$
**Output:** $R, L, U$

$R(i, j) \leftarrow 0, \forall i \in N$ and $\forall j \in P$
Construct the preference graph $G_R$
**while** $\exists$ cycle in $G_R$ **do**
  Eliminate the cycle
  Update $\overline{P}_i$-s by removing any completely assigned paper
  Update $G_R$
**end while**
$U \leftarrow \{i \in N : \overline{P}_i \neq \emptyset\}$
$L \leftarrow$ the last $k_p - |U| + 1$ agents in $N \setminus U$ to have all their submissions completely assigned

*additive preferences are again one example out of a broad class of preference extensions that satisfy consistency.*

**Core.** Our goal is to find a group-fair assignment which treats every possible group of agents at least as well as they could be on their own, thus ensuring that no subset of agents has an incentive to deviate and set up their own separate conference. Formally:

**Definition 23** (Core)**.** *An assignment $R$ is in the core if there is no $N' \subseteq N$, $P'_i \subseteq P_i$ for each $i \in N'$, and restricted assignment $\widehat{R}$ over $N'$ and $P' = \cup_{i \in N'} P'_i$ such that $\widehat{R} \succ_i R$ for each $i \in N'$.*

In words, if any subset of agents deviate with any subset of their submissions and implement any restricted reviewing assignment, at least one deviating agent would not be strictly better off, thus eliminating the incentive for such a deviation. We also remark that our algorithm takes only the preference rankings over individual reviewers $\vec{\sigma}$ as input and produces an assignment $R$ that is guaranteed to be in the core according to *every preference extension* of $\vec{\sigma}$ satisfying order separability and consistency.

## 10.3 CoBRA: An Algorithm for Computing Core-Based Reviewer Assignment

In this section, we prove our main result: when agent preferences are order separable and consistent, an assignment in the core always exists and can be found in polynomial time.

**Algorithm 17** Filling-Gaps

**Input:** $N, P, \vec{\sigma}, k_a, k_p, R, L, U$
**Output:** $R$

**Phase 1**
Construct the greedy graph $G_R$
**while** $\exists$ cycle **do**
  Eliminate the cycle
  Update $\overline{P}_i$-s by removing any completely assigned paper
  Update $U$ and $L$ by moving any agent $i$ from $U$ to $L$ if $\overline{P}_i = \emptyset$
  Update $G_R$
**end while**
**Phase 2**
Construct the topological order $\vec{\rho}$ of $G_R$
**for** $t \in [|U|]$ **do**
  **while** $\overline{P}_{\rho(t)} \neq \emptyset$ **do**
    Pick arbitrary $p_{\rho(t),\ell} \in \overline{P}_{\rho(t)}$
    Find completely assigned $p_{i',\ell'}$, with $R(\rho(t), p_{i',\ell'}) = 0$, for some $i' \in U \cup L \setminus \{\rho(t)\}$
    Find $i'' \neq \rho(t)$ such that $R(i'', p_{\rho(t),\ell}) = 0$ and $R(i'', p_{i',\ell'}) = 1$
    $R(i'', p_{\rho(t),\ell}) \leftarrow 1$; $R(i'', p_{i',\ell'}) \leftarrow 0$; $R(\rho(t), p_{i',\ell'}) \leftarrow 1$
    Remove $p_{\rho(t),\ell}$ from $\overline{P}_{\rho(t)}$ if it is completely assigned
  **end while**
**end for**

**Techniques and key challenges:** The main algorithm CoBRA (Core-Based Reviewer Assignment), presented as Algorithm 15, uses two other algorithms, PRA-TTC and Filling-Gaps, presented as Algorithm 16 and Algorithm 17, respectively. We remark that PRA-TTC is an adaptation of the popular Top-Trading-Cycles (TTC) mechanism, which is known to produce an assignment in the core for the house reallocation problem (and its variants) Shapley and Scarf [1974]. The adaptation mainly incorporates the constraints related to how many papers each reviewer can review and how many reviewers should review each paper. While for $k_p = k_a = 1$, PRA-TTC is identical with the classic TTC that is used for the house reallocation problem, the main difference of this problem with the review assignment problem is that in the latter each agent should give away her item (i.e. her submission) and obtain the item of another agent. Therefore, by simply executing TTC in the review assignment problem, one can get into a deadlock before producing a valid assignment. For example, consider the case of three agents, each with one submission. Each submission must receive one review ($k_p = 1$) and each agent provides one review ($k_a = 1$). The TTC mechanism may start by assigning agents 1 and 2 to review each other's submission, but this cannot be extended into a valid assignment because there is no one left to review the submission of agent 3. This is where Filling-Gaps comes in; it makes careful edits to the partial assignment produced by the PRA-TTC, and the key difficulty is to prove that this produces a valid assignment while still satisfying the core.

| Reviewers \ Rounds | 1 | 2 | 3 | 4 | 5 |
|---|---|---|---|---|---|
| 1 : | $p_3$ | $p_2$ | $p_4$ | | |
| 2 : | $p_1$ | $p_3$ | | $p_6$ | |
| 3 : | $p_2$ | $p_1$ | | | $p_4$ |
| 4 : | | | $p_1$ | | $p_5$ |
| 5 : | | | | $p_2$ | $p_3$ |
| 6 : | | | | $p_5$ | |

(a) Execution of PeerReview-TTC

| Reviewers \ Rounds | 1 | 1 | 2 |
|---|---|---|---|
| 1 : | $p_3, p_2, p_4$ | ~~$p_3$~~, $p_6$ $p_2, p_4$ | $p_6$ $p_2$, ~~$p_4$~~, $p_5$ |
| 2 : | $p_1, p_3, p_6$ | $p_1, p_3, p_6$ | $p_1, p_3, p_6$ |
| 3 : | $p_2, p_1, p_4$ | $p_2, p_1, p_4$ | $p_2, p_1, p_4$ |
| 4 : | $p_1, p_5, p_6$ | $p_1, p_5, p_6$ | $p_1, p_5, p_6$ |
| 5 : | $p_2, p_3$ | $p_2, p_3$ | $p_2, p_3, p_4$ |
| 6 : | $p_5, p_4$ | $p_5, p_4, p_3$ | $p_5, p_4, p_3$ |
| | Phase 1 | Phase 2 | |

(b) Execution of Filling-Gaps

Figure 10.1: Execution of PRCore when $n = 6$, $k_p = k_a = 1$, $\sigma_1 = 2 \succ 3 \succ 4 \succ \ldots$, $\sigma_2 = 3 \succ 1 \succ 5 \succ \ldots$, $\sigma_3 = 1 \succ 2 \succ 5 \succ \ldots$, $\sigma_4 = 1 \succ 3 \succ 5 \succ \ldots$, $\sigma_5 = 6 \succ 4 \succ \ldots$ and $\sigma_6 = 2 \succ \ldots$. On the left table, we see the assignments that are established in each round of PeerReview-TTC by eliminating cycles. After the execution of PeerReview-TTC, three papers, $p_4$, $p_5$, $p_6$ are not completely assigned. Thus, $U = \{4, 5, 6\}$ and $L = \{3\}$. On the right table, we see that there is a cycle in the greedy graph which is eliminated at the first round of Phase 1. In Phase 2, where $\vec{\rho} = (6, 5)$, at the first round, since $p_3$ is authored by an agent in $U \cup L \setminus \{6\}$, is not reviewed by 6 and is completely assigned, $p_3$ is assigned to 6 while it is removed form 1 in which $p_6$ is now assigned. At the second round, since $p_4$ is authored by an agent in $U \cup L \setminus \{5\}$, is not reviewed by 5 and is completely assigned, $p_4$ is assigned to 5 while it is removed form 1 in which $p_5$ is now assigned.

### 10.3.1 Description of CoBRA

Before we describe CoBRA in detail, let us introduce some more notation. Let $m^* = \max_{i \in N} m_i$. For reasons that will become clear later, we want to ensure that $m_i = m^*$, for each $i \in N$. To achieve that, we add $m^* - m_i$ dummy submissions to agent $i$, and the rankings over reviewers with respect to these submissions are arbitrarily. An assignment is called *partial* if there are submissions that are reviewed by less than $k_p$ agents. A submission that is reviewed by $k_p$ agents under a partial assignment is called *completely assigned*. Otherwise, it is called *incompletely assigned*. We denote with $\overline{P}_i(\widehat{R})$ the set of submissions of $i$ that are incompletely assigned under a partial assignment $\widehat{R}$. We omit $\widehat{R}$ from the notation when it is clear from context.

CoBRA calls PRA-TTC, and then if needed, it calls Filling-Gaps. Below, we describe the algorithms.

**PRA-TTC.** In order to define PRA-TTC, we first need to introduce the notion of a *preference graph*. Suppose we have a partial assignment $\widehat{R}$. Each agent $i$ with $\overline{P}_i \neq \emptyset$ picks one of her incompletely assigned submissions arbitrarily. Without loss of generality, we assume that she picks her $\ell^*$-th submission. We define the directed preference graph $G_{\widehat{R}} = (N, E_{\widehat{R}})$ where each agent is a node and for each $i$ with $\overline{P}_i \neq \emptyset$, $(i, i') \in E_{\widehat{R}}$ if and only if $i'$ is ranked highest in $\sigma_{i,\ell^*}$ among the agents that don't review $p_{i,\ell^*}$ and review less than $k_a$ submissions. Moreover, for each $i \in N$ with $\overline{P}_i = \emptyset$, we add an edge from $i$ to $i'$, where $i'$ is an arbitrary agent with $\overline{P}_{i'} \neq \emptyset$. PRA-TTC starts with an empty assignment, constructs the preference graph and searches for a directed cycle. If such a cycle exists, the algorithm eliminates it as following: For each $(i, i')$ that is included in the cycle, it assigns submission $p_{i,\ell^*}$ to $i'$ (if $i$'s submissions are already completely assigned, it does nothing) and removes $p_{i,\ell^*}$ from $\overline{P}_i$, if it is now completely assigned. Then, the algorithm updates the preference graph and continues to eliminate cycles in the same way. When there are no left cycles in the preference graph, the algorithm terminates and returns a set $U$ that contains all the agents that some of their submissions are incompletely assigned and a set $L$ that contains the last $k_p - |U| + 1$ agents whose all submissions became completely assigned.

**Filling-Gaps.** CoBRA calls Filling-Gaps, if the $U$ that returned from PRA-TTC is non empty. Before we describe the Filling-Gaps, we also need to introduce the notion of a *greedy graph*. Suppose that we have a partial assignment $\widehat{R}$ which indicates a set $U$ that contains all the agents whose at least one submission is incompletely assigned. We define the directed greedy graph $G_{\widehat{R}} = (U, E_{\widehat{R}})$ where $(i, i') \in E_{\widehat{R}}$ if $\widehat{R}(i', p_{i,\ell}) = 0$ for some $p_{i,\ell} \in \overline{P}_i$. In other words, while in the preference graph, agent $i$ points only to her favourite potential reviewer with respect to one of her incomplete submissions, in the greedy graph agent $i$ points to any agent in $U \setminus \{i\}$ that could review at least one of her submissions that is incompletely assigned. Filling-Gaps consists of two phases. In the first phase, starting from the partial assignment $R$ that was created from PRA-TTC, it constructs the greedy graph, searches for cycles and eliminates a cycle by assigning $p_{i,\ell}$ to agent $i'$ for each $(i, i')$ in the cycle that exists due to $p_{i,\ell}$ (when an edge exists due to multiple submissions, the algorithm chooses one of them arbitrary). Then, it updates $\overline{P}_i$ by removing any $p_{i,\ell}$ that became completely assigned and also updates $U$ by moving any $i$ to $L$ if $\overline{P}_i$ became empty. It continues by updating the greedy graph and eliminating cycles in the same way. When no more cycles exist in the greedy graph, the algorithm proceeds to the second phase, where in $|U|$ rounds ensures that the incomplete submissions of each agent become completely assigned. Figure 10.1 shows the execution of PRCore in a small instance.

### 10.3.2 Main Result

We are now ready to present our main result.

**Theorem 55.** *When agent preferences are order separable and consistent, each author serves as a reviewer, and $|P_i| \leqslant \lfloor k_a/k_p \rfloor$ for each $i \in N$, CoBRA returns an assignment in the core in polynomial time.*

*Proof.* First, in the next lemma, we show that CoBRA returns a valid assignment.

**Lemma 31.** *CoBRA returns a valid assignment.*

*Proof.* We start by showing that during the execution of PRA-TTC and the execution the first phase of Filling-Gaps, it holds the following lemma.

**Lemma 32.** *During the execution of PRA-TTC and the execution of the first phase of Filling-Gaps, for each $i \in N$ with $\overline{P}_i \neq \emptyset$, it holds that*

$$|R_i^a| = \sum_{j \in [m^*]} |R^p_{p_{i,j}}|.$$

*Proof.* Note that during the execution of PRA-TTC, if an agent $i$ with $\overline{P}_i \neq \emptyset$ is assigned one submission due to the elimination of a cycle, then we know that one of her submissions that is incompletely assigned is also assigned to an agent that does not review it already. Hence, we see that until $\overline{P}_i$ becomes empty, we have that $|R_i^a| = \sum_{j \in [m^*]} |R^p_{p_{i,j}}|$.

Next, we focus on the execution of Filling-Gaps. We know that any $i \in N$ with $\overline{P}_i \neq \emptyset$ is included in $U$. In the first phase, the algorithm eliminates cycles in the greedy graph. With similar arguments as in the elimination of cycles in the preference graph, we get that during and after the first phase of Filling-Gaps,it is still true that $|R_i^a| = \sum_{j \in [m^*]} |R^p_{p_{i,j}}|$ for any $i \in U$ with $\overline{P}_i \neq \emptyset$.

□

Next, we need to show that Line 9 of PRA-TTC is valid which is true if $|U| \leqslant k_p$.

**Lemma 33.** *PRA-TTC returns $|U| \leqslant k_p$.*

*Proof.* For each $i \in U$, from Lemma 32, we know that

$$|R_i^a| = \sum_{j \in [m^*]} |R^p_{p_{i,j}}| < m^* \cdot k_p \leqslant k_a$$

where the first inequality follows since there exists at least one submission of $i$ that is assigned to less than $k_p$ agents. Hence, we get that each $i \in U$ can review more submissions, when PRA-TTC terminates. Now, suppose for contradiction that at the last iteration of PRA-TTC, each agent $i \in U$ has an outgoing edge in the preference graph. In this case, we claim that there exists a directed cycle in the preference graph which is a contradiction since PRA-TTC would not have been terminated yet. To see that, note that each outgoing edge of an agent $i \in U$ either goes to another agent $i' \in U$, since $i'$ can review more submissions, or goes to an agent $i' \notin U$ whose all submissions are completely assigned. In the latter case, $i'$ has an outgoing edge to an agent in $U$ by the definition of the preference graph. Thus, starting

from any agent in $U$, we conclude in an agent in $U$ and eventually we would found a cycle. Therefore, we have that there exists an agent $i^* \in U$ that at the last iteration of the algorithm arbitrary picks her incomplete submission $p_{i^*,\ell^*}$ and does not have any outgoing edge to any other agent. This means that all the agents that can review more submissions, already review $p_{i^*,\ell^*}$. Since all the agents in $U \setminus \{i^*\}$ can review more submissions, we get that all of them are assigned $p_{i^*,\ell^*}$. But since $p_{i^*,\ell^*}$ is not completely assigned, we conclude that $|U \setminus \{i^*\}| < k_p$, which means that $|U| \leqslant k_p$. □

We proceed by showing that Lines 14- 15 of Filling-Gaps are valid.

**Lemma 34.** *When Fillings-Gaps enter the second phase with a non empty $U$, for each $t \in [|U|]$ and for each $p_{\rho(t),\ell} \in \overline{P}_{\rho(t)}$, it exists a completely assigned submission $p_{i',\ell'}$ with $i' \in U \cup L \setminus \{\rho(t)\}$ that is not reviewed by $\rho(t)$, and it also exists an $i'' \in N$ that reviews $p_{i',\ell'}$, but she does not review $p_{\rho(t),\ell}$.*

*Proof.* When $U$ is non empty and no more cycles exists in the greedy graph, the algorithm constructs the topological order of the greedy graph, denoted by $\vec{\rho}$.

First, we show the following proposition.

**Proposition 11.** *For each $t \in [|U|]$, $\rho(t)$ reviews all the incompletely assigned submissions of each $i \in U \setminus \{\rho(1), \ldots, \rho(t-1), \rho(t)\}$.*

*Proof.* Since $\vec{\rho}$ is the topological ordering of the greedy graph, we have that no $i \in U \setminus \{\rho(1), \ldots, \rho(t-1), \rho(t)\}$ has an outgoing edge to $\rho(t)$. But from Lemma 32, we get that $\rho(t)$ can review more submissions, since $\rho(t)$ has submissions that are incompletely assigned which means that

$$\sum_{\ell \in [m^*]} |R^p_{p_{\rho(1),\ell}}| < kp \cdot m^* \leqslant k_a.$$

Therefore, from the definition of the greedy graph, we get that $\rho(t)$ reviews all the incompletely assigned submissions of each $i$ in $U \setminus \{\rho(1), \ldots, \rho(t-1), \rho(t)\}$. □

Next, we show by induction that for each $t \in [|U|]$ as long as $\overline{P}_{\rho(t)}$ is non-empty and it holds that $|R^a_{\rho(t)}| = \sum_{\ell \in [m^*]} |R^p_{p_{\rho(t),\ell}}|$, there exists a completely assigned submission $p_{i',\ell'}$ of an agent $i' \in U \cup L \setminus \{\rho(t)\}$ that is not reviewed by $\rho(t)$, and $i'' \in N$ that reviews $p_{i',\ell'}$ and does not review $p_{\rho(t),\ell} \in \overline{P}_{\rho(t)}$.

We start with $t = 1$. First, suppose for contradiction that $\rho(1)$ reviews all the submissions of all the agents in $U \cup L \setminus \{\rho(1)\}$. Then, we would have that

$$|R^a_{\rho(1)}| = \sum_{\ell \in [m^*]} |R^p_{p_{\rho(1),\ell}}| = k_p \cdot m^*,$$

where the first equality follows from the assumption that $|R^a_{\rho(1)}| = \sum_{\ell \in [m^*]} |R^p_{p_{\rho(1),\ell}}|$ and the second inequality follows from the facts that $|U \cup L \setminus \{\rho(1)\}| = k_p$ and each agent has $m^*$ submissions. But then we would conclude that all the submissions of $\rho(1)$ are completely assigned since $\rho(1)$ has $m^*$ submissions and each of them should be assigned to $k_p$ reviewers which is a contradiction. Moreover, from Proposition 11 we know that $\rho(1)$ reviews all the incomplete assigned submissions that belongs to some agent $i \in U \setminus \{\rho(1)\}$. Hence, we get that since $\rho(1)$ reviews all the incompletely assigned submissions but she cannot review all the submissions of all agents in $i \in U \cup L \setminus \{\rho(1)\}$, there exists

a completely assigned submission $p_{i',\ell'}$ that belongs to some $i' \in U \cup L \setminus \{\rho(1)\}$ and is not reviewed by $\rho(1)$. In addition, since $p_{i',\ell'}$ is reviewed by $k_p$ agents and not from $\rho(1)$, while $p_{\rho(1),\ell} \in \overline{P}_{\rho(1)}$ is reviewed by strictly less than $k_p$ agents, it exists an agent $i''$ that reviews the former submission but not the latter. It remains to show that during the execution of the 1-th iteration of the for loop in the second phase of Filling-Gaps, it holds that $|R^a_{\rho(1)}| = \sum_{\ell \in [m^*]} |R^p_{p_{\rho(1),\ell}}|$. Note that if every time that the algorithm enters the second while loop of the algorithm, this property is satisfied, then the property remains true at the end of this execution, since as we show above, in this case there are $p_{i',\ell'}$ and $i''$ with the desired properties, and therefore one incompletely assigned submission of $\rho(1)$ is assigned to a new reviewer and concurrently $\rho(1)$ is assigned a new submission to review. We get that $|R^a_{\rho(1)}| = \sum_{\ell \in [m^*]} |R^p_{p_{\rho(1),\ell}}|$ is true during the execution of the 1-st ieration of the for loop by noticing that from Lemma 32, we know that this is true when we first enter the while loop of the second phase.

Suppose that the hypothesis holds for $t-1$. Note that from the base case and the hypothesis, at iteration $t$, all the submissions of each agent in $i' \in L \cup \{\rho(1), \ldots, \rho(t-1)\}$ are completely assigned. Thus, any incompletely assigned submission, that does not belong to $\rho(t)$, belongs to some agent $i \in U \setminus \{\rho(1), \ldots, \rho(t-1), \rho(t)\}$. But, from Proposition 11 we already know that $\rho(t)$ reviews any such submission. Moreover, we note that $\rho(t)$ cannot review all the submissions of all the agents in $U \cup L \setminus \{\rho(t)\}$. Indeed, if we assume for contradiction that $\rho(t)$ reviews all the submissions of all the agents in $U \cup L \setminus \{\rho(t)\}$, then we have that

$$|R^a_{\rho(t)}| = \sum_{\ell \in [m^*]} |R^p_{p_{\rho(t),\ell}}| = k_p \cdot m^*,$$

where the first inequality follows from the assumption that $|R^a_{\rho(t)}| = \sum_{\ell \in [m^*]} |R^p_{p_{\rho(t),\ell}}|$ and the second follows from the facts that $|U \cup L \setminus \{\rho(t)\}| = k_p$ and each of them has $m^*$ submissions, which would imply that all the submissions of $\rho(t)$ are completely assigned. Hence, we get that since $\rho(t)$ reviews all the incompletely assigned submissions but cannot review all the submissions of all agents in $i \in U \cup L \setminus \{\rho(t)\}$, there exists a completely assigned submission that belongs to some $i' \in U \cup L \setminus \{\rho(t)\}$ and is not reviewed by $\rho(t)$. Moreover, we show that there exists $i''$ that reviews $p_{i',\ell'}$, but does not review $p_{\rho(t),\ell} \in \overline{P}_{\rho(t)}$. Indeed, since $p_{i',\ell'}$ is reviewed by $k_p$ agents and not from $\rho(t)$, while $p_{\rho(t),\ell}$ is reviewed by strictly less than $k_p$ agents, it exists an agent that reviews the former submission but not the latter. It remains to show that during the execution of the $t$-th iteration of the for loop in the second phase of Filling-Gaps, it holds that $|R^a_{\rho(t)}| = \sum_{\ell \in [m^*]} |R^p_{p_{\rho(t),\ell}}|$. Note that if when we first enter the while loop of the $t$-th iteration it is indeed true that $|R^a_{\rho(t)}| = \sum_{\ell \in [m^*]} |R^p_{p_{\rho(t),\ell}}|$, then during the whole execution of the while loop this property remains true, since as we show above, in this case there are $p_{i',\ell'}$ and $i''$ with the desired properties. From Lemma 32, we know that this is true when we enter the second phase of Filling-Gaps. Before, round $t$, if $\rho(t)$ is assigned a new submission to review, she is removed one of the old assigned submissions, while none of her incompletely assigned submissions is assigned to any agent. Hence, indeed we have the desired property, when we first enter the $t$-th round. □

We partition the agents in $N$, into two parts $N_1$ and $N_2$, where $N_1$ contains all the agents that do not belong in $U$ that PRA-TTC returns and $N_2 = N \setminus N_1$. We proceed by separately showing that the assignment that CoBRA returns is valid over the agents in $N_1$ and over the agent in $N_2$.

**Valid Assignment over the agents in $N_1$.** Note that in PRA-TTC, each submission is assigned to at most $k_p$ reviewers and therefore, during the execution of PRA-TTC, for each $i \in N$, it holds that

$\sum_{j\in[m^*]} |R^p_{p_{i,j}}| \leqslant k_p \cdot m^*$. From Lemma 32, since $k_p \cdot m \leqslant k_a$, we get that in PRA-TTC, each $i \in N_1$ is not assigned more than $k_a$ papers to review until the point where all of her submissions become completely assigned. After that point, an agent may still participate in a cycle as long as she reviews strictly less than $k_a$ submissions. Therefore, when we exit PRA-TTC, each agent in $N_1$ does not review more than $k_a$ submissions and all her submissions are completely assigned. In Filling-Gaps, from Lemma 34, we get that if an agent in $N_1$ is assigned a new submission to review, she is removed one of the submissions that she already reviews. Moreover, the assignments of submissions that belong to agents in $N_1$ do not change. Hence, we can conclude that the assignment that CoBRA returns is valid with respect to the agents in $N_1$.

**Valid Assignment over the agents in $N_2$.** From Lemma 32, we get that during the execution of PRA-TTC each agent $i$ that is included in $U$ that PRA-TTC returns reviews less than $k_a$ submissions, since some of her submissions are not completely assigned (which means that $\sum_{\ell\in[m^*]} |R^p_{p_{i,\ell}}| < k_p \cdot m^*$). From the same lemma, we have that after the execution of the first phase of Filling-Gaps, it holds that $|R^a_i| = \sum_{\ell\in[m^*]} |R^p_{p_{i,\ell}}|$. Next, we show that this property remains true after the second phase of Filling-Gaps. Indeed, from Lemma 34, we have that during the second phase of Filling-Gaps, if $i \in U$ is assigned a new submission to review without one of her incompletely assigned submissions is assigned to a new reviewer, she is removed one of her assigned submissions; on the other hand, if one of her incompletely assigned submissions is assigned to a new reviewer, she is also assigned to review a new submission. Lastly, from Lemma 34, we conclude that in the second phase of Filling-Gaps, all the submissions become eventually completely assigned since in each iteration of the while loop, an incompletely assigned submission is assigned to one more reviewer. Therefore in the assignment that Filling-Gaps returns, no agent in $N_2$ reviews more than $k_a$ submissions and all the submissions of the agents in $N_2$ are completely assigned, which means that the assignment is valid with respect to the agents in $N_2$ as well. □

Now, we show that the final assignment $R$ that CoBRA returns is in the core. Note that while it is possible that an assignment of a submission of an agent in $U \cup L$, that was established during the execution of PRA-TTC, to be removed in the execution of Filling-Gaps, this never happens for submissions that belong to some agent in $N \setminus (U \cup L)$. For the sake of contradiction, assume that $N' \subseteq N$, with $P'_i \subseteq P_i$ for each $i \in N'$, deviate to a restricted assignment $\widetilde{R}$ over $N'$ and $\cup_{i\in N'} P'_i$ . Note that $\widetilde{R}$ is valid only if $|N'| > k_p$, as otherwise there is no way each submission in $\cup_{i\in N'} P'_i$ to be completely assigned, since no agent can review her own submissions.

We distinguish into two cases and we show that in both cases the assignment is in the core.

**Case I: $\exists i \in N' : i \notin L \cup U$.** Let $i^* \in N'$ be the first agent in $N'$ whose all submissions became completely assigned in the execution of PRA-TTC. Note that since there exists $i \notin U \cup L$, we get that $i^* \notin U \cup L$ from the definitions of $U$ and $L$. Now, consider any $p_{i^*,\ell}$. Let $Q_1 = R^p_{p_{i^*,\ell}} \setminus (R^p_{p_{i^*,\ell}} \cap \widetilde{R}^p_{p_{i^*,\ell}})$ and $Q_2 = \widetilde{R}^p_{p_{i^*,\ell}} \setminus (R^p_{p_{i^*,\ell}} \cap \widetilde{R}^p_{p_{i^*,\ell}})$. If $Q_1 = \emptyset$, then we have that $R^p_{p_{i^*,\ell}} = \widetilde{R}^p_{p_{i^*,\ell}}$ which means that $R^p_{p_{i^*,\ell}} \succeq_{i^*,\ell} \widetilde{R}^p_{p_{i^*,\ell}}$. Otherwise, let $i' = \arg\max_{i\in Q_1} \sigma_{i^*,\ell}(i)$, i.e. $i'$ is ranked at the lowest position in $\sigma_{i^*,\ell}$ among the agents that review $p_{i^*,\ell}$ under $R$ but not under $\widetilde{R}$. Moreover, let $i'' = \arg\min_{i\in Q_2} \sigma_{i^*,\ell}(i)$, i.e. $i''$ is ranked at the highest position in $\sigma_{i^*,\ell}$ among the agents that review $p_{i^*,\ell}$ under $\widetilde{R}$ but not under $R$. We have $R(i', p_{i^*,\ell}) = 1$, if and only if $i^*$ has an outgoing edge to $i'$ at some round of PRA-TTC. At the same round, we get that $i''$ can review more submissions, since $i'' \in N'$ and if $i^*$

has incompletely assigned submissions, then any $i \in N'$ has incompletely assigned submissions, and hence $|R^a_{i''}| < k_p \cdot m^* \leqslant k_a$. This means that if $\sigma_{i^*,\ell}(i') > \sigma_{i^*,\ell}(i'')$, then $i^*$ would point $i''$ instead of $i'$. We conclude that $\sigma_{i^*,\ell}(i') < \sigma_{i^*,\ell}(i'')$. Then, from the definition of $i'$ and $i''$ and from the order separability property we have that $R^p_{p_{i^*,\ell}} \succ_{i^*,\ell} \widetilde{R}^p_{p_{i^*,\ell}}$. Thus, either if $Q_1$ is empty or not, we have that for any $p_{i^*,\ell} \in P'_i$, it holds that $R^p_{p_{i^*,\ell}} \succeq_{i^*,\ell} \widetilde{R}^p_{p_{i^*,\ell}}$ and from consistency, we get that $R \succeq_{i^*} \widetilde{R}$ which is a contradiction.

**Case II: $\nexists i \in N' : i \notin L \cup U$.** In this case we have that $N' = U \cup L$, as $|U \cup L| = k_p + 1$. This means that for each $i \in U \cup L$ and $\ell \in [m^*]$, $\widetilde{R}^p_{p_{i,\ell}} = (U \cup L) \setminus \{i\}$. Let $i^* \in L$ be the first agent in $L$ whose all submissions became completely assigned in the execution of PRA-TTC. Consider any $p_{i^*,\ell}$. Note that it is probable that while $p_{i^*,\ell}$ was assigned to some agent $i$ in PRA-TTC, it was moved to another agent $i'$ during the execution of Filling-Gaps. But, $i'$ belongs to $U$ and we can conclude that if $p_{i^*,\ell}$ is assigned to some $i \in N \setminus U$ at the output of CoBRA, this assignment took place during the execution of PRA-TTC. Now, let $Q_1 = R^p_{p_{i^*,\ell}} \setminus (R^p_{p_{i^*,\ell}} \cap \widetilde{R}^p_{p_{i^*,\ell}})$ and $Q_2 = \widetilde{R}^p_{p_{i^*,\ell}} \setminus (R^p_{p_{i^*,\ell}} \cap \widetilde{R}^p_{p_{i^*,\ell}})$. If $Q_1 = \emptyset$, then we have that $R^p_{p_{i^*,\ell}} = \widetilde{R}^p_{p_{i^*,\ell}}$ which means that $R^p_{p_{i^*,\ell}} \succeq_{i^*,\ell} \widetilde{R}^p_{p_{i^*,\ell}}$. If $Q_1 \neq \emptyset$, then $Q_1 \subseteq N \setminus (U \cup L)$ and $Q_2 \subseteq U \cup L$ since $\widetilde{R}^p_{p_{i^*,\ell}} = U \cup L$. Let $i' = \arg\max_{i \in Q_1} \sigma_{i^*,\ell}(i)$, i.e. $i'$ is ranked at the lowest position in $\sigma_{i^*,\ell}$ among the agents that review $p_{i^*,\ell}$ under $R$ but not under $\widetilde{R}$. Moreover, let $i'' = \arg\min_{i \in Q_2} \sigma_{i^*,\ell}(i)$, i.e. $i''$ is ranked at the highest position in $\sigma_{i^*,\ell}$ among the agents that review $p_{i^*,\ell}$ under $\widetilde{R}$ but not under $R$. From above, we know that the assignment of $p_{i^*,\ell}$ to $i'$ was implemented during the execution of PRA-TTC, since $i' \in N \setminus (U \cup L)$. Hence, with very similar arguments as in the previous case, we will conclude that $\sigma_{i^*,\ell}(i') < \sigma_{i^*,\ell}(i'')$. We have $R(i', p_{i^*,\ell}) = 1$ if and only if $i^*$ has an outgoing edge to $i'$ at some round of PRA-TTC. At this round, we know that $i''$ can review more submissions, since $i'' \in N'$ and if $i^*$ has incompletely assigned submissions, then any $i \in N'$ has incompletely assigned submissions. This means that if $\sigma_{i^*,\ell}(i') > \sigma_{i^*,\ell}(i'')$, then $i^*$ would point $i''$ instead of $i'$. Hence, we conclude that $\sigma_{i^*,\ell}(i') < \sigma_{i^*,\ell}(i'')$. Then, from the definition of $i'$ and $i''$ and from the order separability property we have that $R^p_{p_{i^*,\ell}} \succ_{i^*,\ell} \widetilde{R}^p_{p_{i^*,\ell}}$. Thus, either if $Q_1$ is empty or not, we have that for any $p_{i^*,\ell} \in P'_i$, it holds that $R^p_{p_{i^*,\ell}} \succeq_{i^*,\ell} \widetilde{R}^p_{p_{i^*,\ell}}$ and from consistency we get that $R \succeq_{i^*} \widetilde{R}$ which is a contradiction. □

## 10.4 The Core and Other Objective

In the previous section, we show that an assignment in the core is guaranteed to exist under very minor assumptions regarding the preferences of the authors. As we mentioned in the introduction, existing works have focused on different objectives. To be able to compare our objective with the existing ones, from now on, we take the standard approach that for each paper $j$ and each reviewer $i$, it is given a similarity score $S(i, j)$ which is calculated as a function of different parameters. Given an assignment $R$, we also assume that the utilities of the papers and the authors are additive, i.e. the utility of a paper $j$, denoted by $u^p_j$, is equal to $u^p_j = \sum_{i \in R^p_j} S(i, j)$ and the utility of an author $i$, denoted by $u^a_i$, is equal to $u^a_i = \sum_{j \in P_i} u^p_j$.

The most known objective, as it is used at the Toronto Paper Matching System (TPMS) [Charlin and Zemel, 2013] is the maximization of the utilitarian social welfare (USW), which is given by

$$USW(R) = \sum_{i \in N} \sum_{j \in P_i} u^p_j(R).$$

We denote the algorithm which computes such an assignment as TPMS.

A different objective that was introduced by Stelmakh et al. [2019] is to maximize the egalitarian social welfare (ESW) which is given by

$$ESW(R) = \min_{j \in \cup_{i \in N} P_i} u_j^p(R).$$

Stelmakh et al. [2019] considered the extended leximin version of this objective where subject to maximize the minimum utility of all papers, they aim to maximize the second minimum utility of all papers, and subject to that they aim to maximize the third minimum utility of all papers and so on. The algorithm that achieves this objective is called PeerReview4All (PR4A).

A reasonable question is whether the core is compatible with good approximations of USW and ESW. Below, we show that there are instances where any solution in the core does not achieve an approximation ratio better than $\Omega(n)$ with respect to USW and a finite approximation ratio with respect to ESW.

**Theorem 56.** *For any $n$, $k_p$ and $k_a$, where $k_a$ is divisible by $k_p$, when $n \geqslant 2 \cdot k_p \cdot ka + 1$, there exists an instance such that no assignment in the core achieves approximation ratio better than $\Omega(n/k_a)$ with respect to USW and a finite approximation with respect to ESW.*

*Proof.* Suppose that each agent submits $k_a/k_p$ submissions. Let $N_1 = \{1, \ldots, \lfloor n/2 \rfloor\}$ and $N_2 = \{\lfloor n/2 \rfloor + 1, \ldots, n\}$ Consider an instance where the similarity scores are as following:

- For $i, i' \in N_1$, $s(i', p_{i,1}) = 0$
- For $i \in N_1$ and $i' \in N_2$, $s(i', p_{i,1}) = 1$
- For $i, i' \in N_2$, $s(i', p_{i,1}) = \epsilon_1$
- For $i \in N_2$ and $i' \in N_1$, $s(i', p_{i,1}) = \epsilon_2$
- For $i, i' \in N$, $s(i', p_{i,j}) = 0$, for each $j > 1$

where $\epsilon_1 > \epsilon_2$.

Now, suppose that there are at least $k_p + 1$ agents in $N_2$ whose first submissions are not exclusively assigned to reviewers in $N_2$. Then, if they deviate and assign their submissions among themselves, this would lead in a valid assignment as each submission would be reviewed by $k_p$ agents and they would strictly improve their utility. Thus, we conclude that an assignment is in the core if the first submission of at most $k_p$ authors in $N_2$ are not exclusively assigned to authors in $N_2$. Hence, we get that there are at most $k_p \cdot k_a$ assignments of submissions that belong in authors in $N_1$ to reviewers in $N_2$. This means that the maximum utilitarian welfare of an assignment in the core is equal to $k_p \cdot k_a + n \cdot k_a \cdot \epsilon_1$ where as $\epsilon_1$ goes to zero the welfare goes to $k_p \cdot k_a$. Moreover, for $\lfloor n/2 \rfloor > k_p \cdot k_a$, we have that at least one agent in $N_1$ should have zero utility under any assignment in the core. On the other hand, by assigning the submissions of the agents in $N_1$ to agents in $N_2$ and the submissions of the agents in $N_2$ (except for the submissions of the last agent when $n$ is odd) to agents in $N_1$, we achieve utility at least equal to $k_p \cdot \lfloor n/2 \rfloor$. Thus, the approximation with respect to the optimal social welfare cannot be better than $\Omega(n/k_a)$. Moreover, by this way the minimum utility is equal to $\epsilon_2$ and for $\epsilon_2 > 0$, and we get that the approximation with respect to ESW is unbounded. □

The condition of the theorem is quite realistic since in practice $k_p$ and $k_a$ are small constants. Thus, we can also get the following Corollary.

**Corollary 57.** *There are instances where no assignment in the core achieves approximation ratio better than $\Omega(n)$ with respect to USW.*

It is known from the literature that we can find an assignment with maximum USW in polynomial times using standard tools [Taylor, 2008]. In the previous section, we also presented a polynomial time algorithm that finds an assignment in the core. Here, we surprisingly show that if $NP \neq P$, there is no polynomial time algorithm that finds an assignment in the core with maximum USW. Moreover, it is known that finding an assignment with maximum egalitarian social welfare is a NP-hard problem [Demko and Hill, 1988]. Here, we show that it is NP-hard to find an assignment in the core with bounded approximation to the maximum ESW that can be achieved by any assignment in the core. The proof can be found in the appendix.

**Theorem 58.** *Finding an assignment in the core with maximum social welfare is NP-hard. Finding an assignment in the core with bounded approximation with respect to the maximum egalitarian welfare achieved by assignments in the core is NP-hard.*

*Proof.* We begin by proving the theorem for the case that $k_a = k_p = 1$, and later we generalize it for any $k_a$.

We use a polynomial-time reduction from 2P2N-3SAT, the special case of 3SAT where every boolean variable appears twice as positive and twice as negative literal. Let $\phi$ be an instance of 2P2N-3SAT which consists of $n$ boolean variables, $x_1, x_2, ..., x_n$ and $m$ clauses $C_1, C_2, ..., C_m$ with $n = 3m/4$. We assume that $n$ is divisible by 3. Given $\phi$, we construct an instance for the assignment review problem such that:

- If $\phi$ is satisfiable, then there exists an assignment in the core with social welfare at least $4n/3$ and with minimum paper score $\epsilon > 0$.
- If $\phi$ is not satisfiable, then any assignment in the core has social welfare less than $4n/3 - 1/2$ and minimum paper score equal to zero.

The assignment review problem is as following: For each boolean variable $x_i$, we add agents $x_{i,1}$, $x_{i,2}$, $\overline{x}_{i,2}$ and $\overline{x}_{i,1}$, where the first two (respectively, the last two) agents corresponds to the two occurrences of the literal $x_i$ (resp., $\neg x_i$). Moreover, for each clause $C_j$ we add agent $c_j$. Each agent has exactly on submission and the expertises of the agents over the submissions are as follows: For each bolean variable $x_i$,

- $S(x_{i,1}, p_{\overline{x}_{i,1}}) = S(\overline{x}_{i,1}, p_{x_i}) = \epsilon_1$
- $S(x_{i,2}, p_{\overline{x}_{i,2}}) = S(\overline{x}_{i,2}, p_{x_{i,2}}) = \epsilon_1$
- $S(x_{i,1}, p_{\overline{x}_{i,2}}) = S(\overline{x}_{i,2}, p_{x_{i,1}}) = \epsilon_2$
- $S(x_{i,2}, p_{\overline{x}_{i,1}}) = S(\overline{x}_{i,1}, p_{x_{i,2}}) = \epsilon_2$

and for each $C_j = (\ell_1 \vee \ell_2 \vee \ell_3)$ where $\ell_1$, $\ell_2$ and $\ell_3$ correspond to the $t_1$-th appearance of literal $x_{i_1}$ (resp., of the literal $\neg x_{i_1}$), $t_2$-th appearance of literal $x_{i_3}$ (resp., of the literal $\neg x_{i_3}$) and $t_3$-th appearance of literal $x_{i_3}$ (resp., of the literal $\neg x_{i_3}$), respectively, with $t_1, t_2, t_3 \in \{1, 2\}$,

- $S(x_{i_1,t_1}, p_{c_j}) = S(x_{i_2,t_2}, p_{c_j}) = S(x_{i_3,t_3}, p_{c_j}) = 1$
- $S(i, p_{c_j})$=0, for any agent $i \notin \{x_{i_1,t_1}, x_{i_2,t_12}, x_{i_3,t_3}\}$

and all the remaining scores are equal to $\epsilon_3$ with $1/(8n) \geqslant \epsilon_1 > \epsilon_2 > \epsilon_3 > 0$.

If $\phi$ is satisfiable, then we use a truth assignment to find an assignment in the core with social welfare at least $4n/3$ and minimum paper score equal to $\epsilon_2 > 0$ as follows: For every true variable $x_i$, we assign $p_{x_{i,t}}$ to $\overline{x}_{i,t}$, for each $t \in \{1, 2\}$. Respectively, for every false variable $x_i$, we assign $p_{\overline{x}_{i,t}}$ to $x_{i,t}$, for each $t \in \{1, 2\}$. For every clause $C_j$, we arbitrary select one of the true literals of the clause, and assign $p_{c_j}$ to the corresponding agent. All the remaining assignments are arbitrary. First, we see that under this assignment each agent $c_j$ has utility 1, as $p_{c_j}$ is assigned to a reviewer with similarity score equal to 1. Hence, the social welfare is at least $4n/3$. Moreover, the minimum score utility is equal to $\epsilon_2$. Now, we show that it is also in the core. First, no agent $c_j$ has incentives to deviate as her submission is assigned to one of the best possible agents for her submission. Now, consider an agent $x_{i,t}$, for $t \in \{1, 2\}$, when variable $x_i$ is true. From the construction, $p_{x_{i,t}}$ is assigned to $\overline{x}_{i,t}$ which has the highest similarity score for it and hence $x_{i,t}$ does not have any incentives to deviate. With similar, arguments we can show that $\overline{x}_{i,t}$, for $t \in \{1, 2\}$, does not have incentives to deviate when $x_i$ is false. Next, consider an agent $x_{i,t}$, for $t \in \{1, 2\}$ when variable $x_i$ is false. While $\overline{x}_{i,1}$ and $\overline{x}_{i,2}$ are the two agents that have higher similarity score for $p_{x_{i,t}}$, than its current reviewer, we know from above that $\overline{x}_{i,1}$ and $\overline{x}_{i,2}$ do not have any incentives to deviate. Hence, $x_{i,t}$ does not have any incentive to deviate with any other agents when $x_i$ is false, and similarity we can show that $\overline{x}_{i,t}$ does not have any incentive to deviate if $x_i$ is true. Thus, in any case there is no deviating coalition and the assignment is in the core.

Now, we show that if the social welfare is at least $4n/3$, then $\phi$ is satisfiable. Moreover, if the lower utility of each paper is positive, then $\phi$ is satisfiable. First, assume for the sake of contradiction that some $p_{c_j}$ is not assigned to an agent that corresponds to one of the literals of $C_j$. Then, since any other agent has similarity score equal to 0 for $p_{c_j}$, the sum of utilities of agents $c_j$ for $j \in [m]$ is at most $4n/3 - 1$. But since for any $p_{x_{i,t}}$ and $p_{\overline{x}_{i,t}}$ the similarity score of any agent is less than $1/8n$ and there are $4n$ such submissions, we have that the overall sum of utilities is at most $4n/3 - 1/2$ and we reach a contradiction. Moreover, we see that in this case the minimum utility of a paper is equal to zero. Hence, we conclude that every $p_{c_j}$ is assigned to an agent that corresponds to one of the literals of $C_j$. Suppose that $p_{c_j}$ is assigned to $x_{i,t}$, where $x_i$ appears as positive literary to $C_j$, and without loss of generality assume that $t = 1$. Then, we see that $p_{x_{i,1}}$ is assigned to $\overline{x}_{i,1}$, as otherwise $x_{i,1}$ and $\overline{x}_{i,1}$ could deviate to their own coalition. Moreover, we notice that the assignment is in the core if and only if at least one of $x_{i,2}$ and $p_{\overline{x}_{i,2}}$ reviews the other's submission. For the sake of contradiction, assume that the assignment meets this requirement by assigning $p_{\overline{x}_{i,2}}$ to $x_{i,2}$, but not $p_{x_{i,2}}$ to $\overline{x}_{i,2}$. Then, $\overline{x}_{i,1}$ and $x_{i,2}$ as none of them reviews each other's submission, they could deviate to their own coalition. So, it should be the case that $\overline{x}_{i,2}$ reviews $p_{x_{i,2}}$. Thus, we have that if $p_{c_j}$ is assigned to $x_{i,1}$, then $p_{x_{i,1}}$ and $p_{x_{i,2}}$ are assigned to $\overline{x}_{i,1}$ and $\overline{x}_{i,2}$, respectively. Hence, no $p_{c_{j'}}$ can be assigned to $\overline{x}_{i,1}$ or $\overline{x}_{i,2}$. With similar arguments, we can show that if $p_{c_j}$ is assigned to $\overline{x}_{i,t}$, then there is no $p_{c_{j'}}$ that is assigned to $x_{i,1}$ or $x_{i,2}$. Now, we see that by setting the variable $x_i$ to true if some $p_{c_j}$ is reviewed by $x_{i,1}$ or $x_{i,2}$, and to false if some $p_{c_j}$ is reviewed by $\overline{x}_{i,1}$ or $\overline{x}_{i,2}$ (these cannot happen concurrently), we get an assignment that satisfies all clauses of $\phi$.

Now, when $k_a > 1$, then for each agent $x_{i,t}$ (resp. $\overline{x}_{i,t}$) in the above construction, we assume that there are $k_a - 1$ other agents, denoted by $x_{i,t}^1, \ldots, x_{i,t}^{k_a}$ such that each of them have similarity score equal to $\epsilon_4$ for each other's submission with $1/(8n) > \epsilon_4 > \epsilon_1$. The remaining scores are set as for the case of $x_{i,t}$. If there are at least two agents among $x_{i,t}^1, \ldots, x_{i,t}^{k_a}$ that their submissions are not reviewed by another agent in $\{x_{i,t}^1, \ldots, x_{i,t}^{k_a}\}$, then they have incentives to deviate. Hence, we see that at most one of them can review a submission that does not belong to some agent in $\{x_{i,t}^1, \ldots, x_{i,t}^{k_a}\}$ and at most one submission of some agent in $\{x_{i,t}^1, \ldots, x_{i,t}^{k_a}\}$ may not reviewed from some agent in the same set. Thus, by interpreting this agent as $x_{i,t}$ and this submission as $p_{x_{i,t}}$, with arguments as above the statement follows.

□

## 10.5 Experiments

In this section, we empirically compare CoBRA to TPMS [Charlin and Zemel, 2013], which is widely used (for example, it was used by NeurIPS for many years), and PR4A [Stelmakh et al., 2019], which was used in ICML 2020 Stelmakh [2021]. As mentioned in the introduction, these algorithms assume the existence of a similarity or affinity score for each pair of reviewer $i$ and paper $j$, denoted by $S(i, j)$. The score (or utility) of a paper under an assignment $R$, denoted by $u_j^p$, is computed as $u_j^p \triangleq \sum_{i \in R_j^p} S(i, j)$. TMPS finds an assignment $R$ that maximizes the utilitarian social welfare (USW), i.e., the total paper score $\sum_{j \in P} u_j^p$, whereas PR4A finds an assignment that maximizes the egalitarian social welfare (ESW), i.e., the minimum paper score $\min_{j \in P} u_j^p$.[6] We use $k_a = k_p = 3$ in these experiments.[7]

**Datasets.** We use three conference datasets: from the Conference on Computer Vision and Pattern Recognition (CVPR) in 2017 and 2018, which were both used by Kobren et al. [2019], and from the International Conference on Learning Representations (ICLR) in 2018, which was used by Xu et al. [2019]. In the ICLR 2018 dataset, similarity scores and conflicts of interest are also available. While a conflict between a reviewer and a paper does not necessarily indicate authorship, it is the best indication we have available, so, following Xu et al. [2019], we use the conflict information to deduce authorship. Since in our model each submission has one author, and no author can submit more than $\lfloor k_a/k_p \rfloor = 1$ papers, we compute a maximum cardinality matching on the conflict matrix to find author-paper pairs, similarly to what Dhull et al. [2022] did. In this way, we were able to match 883 out of the 911 papers. We disregard any reviewer who does not author any submissions, but note that the addition of more reviewers can only improve the results of our algorithm since these additional reviewers have no incentive to deviate. For the CVPR 2017 and CVPR 2018 datasets, similarity scores was available, but not the conflict information. In both these datasets, there are fewer reviewers than papers. Thus, we constructed artificial authorship relations by sequentially processing papers and matching each paper to the reviewer with the highest score for it, if this reviewer is still unmatched. In this way, we were able to match $1373$ out of $2623$ papers from CVPR 2017 and $2840$ out of $5062$ papers from CVPR 2018. In the ICLR 2018 and CVPR 2017 datasets, the similarity scores take values in $[0, 1]$, so we accordingly normalized the

[6] Technically, subject to this, it maximizes the second minimum paper score, and then the third minimum paper score, etc. This refinement is also known as leximin in the literature.

[7] Note that $k_p = 3$ reviews per submission is quite common, although the reviewer load $k_a$ is typically higher in many conferences, often closer to 6. However, the differences between different algorithms diminish with higher values of $k_a$.

| Dataset | Alg | USW | ESW | $\alpha$-Core | | CV-Pr |
|---|---|---|---|---|---|---|
| | | | | #unb-$\alpha$ | $\alpha^*$ | |
| CVPR '17 | CoBRA | $1.225 \pm 0.021$ | $0.000 \pm 0.000$ | $0\%$ | $1.000 \pm 0.000$ | $0\%$ |
| | TPMS | $1.497 \pm 0.019$ | $0.000 \pm 0.000$ | $89\%$ | $3.134 \pm 0.306$ | $100\%$ |
| | PR4A | $1.416 \pm 0.019$ | $0.120 \pm 0.032$ | $51\%$ | $1.700 \pm 0.078$ | $100\%$ |
| CVPR '18 | CoBRA | $0.224 \pm 0.004$ | $0.004 \pm 0.001$ | $0\%$ | $1.000 \pm 0.000$ | $0\%$ |
| | TPMS | $0.286 \pm 0.005$ | $0.043 \pm 0.004$ | $0\%$ | $1.271 \pm 0.038$ | $100\%$ |
| | PR4A | $0.282 \pm 0.005$ | $0.099 \pm 0.001$ | $0\%$ | $1.139 \pm 0.011$ | $100\%$ |
| ICLR '18 | CoBRA | $0.166 \pm 0.001$ | $0.028 \pm 0.001$ | $0\%$ | $1.000 \pm 0.000$ | $0\%$ |
| | TPMS | $0.184 \pm 0.001$ | $0.048 \pm 0.002$ | $0\%$ | $1.048 \pm 0.008$ | $90\%$ |
| | PR4A | $0.179 \pm 0.001$ | $0.082 \pm 0.001$ | $0\%$ | $1.087 \pm 0.009$ | $100\%$ |

Table 10.1: Results on CVPR 2017 and 2018, and ICLR 2018.

CVPR 2018 scores as well.

**Measures.** We are most interested in measuring the extent to which the existing algorithms provide incentives for communities of researchers to deviate. To quantify this, we need to specify the utilities of the authors. We assume that they are additive, i.e., the utility of each author in an assignment is the total similarity score of the $k_p = 3$ reviewers assigned to their submission.

*Core violation factor:* Following the literature Fain et al. [2018], we measure the multiplicative violation of the core (if any) that is incurred by TPMS and PR4A (CoBRA provably does not incur any). This is done by computing the maximum value of $\alpha \geqslant 1$ for which there exists a subset of authors such that by deviating and implementing some valid reviewing assignment of their papers among themselves, they can each improve their utility by a factor of at least $\alpha$. This can easily be formulated as a binary integer linear program (BILP). Because this optimization is computationally expensive (the most time-consuming component of our experiments), we subsample $100$ papers from each dataset in each run, and report results averaged over $100$ runs along with $95\%$ confidence intervals in Table 10.1.

Note that whenever there exists a subset of authors with zero utility each in the current assignment who can deviate and receive a positive utility each, the core deviation $\alpha$ becomes infinite. We separately measure the percentage of runs in which this happens (in the column #unb-$\alpha$), and report the average $\alpha$ among the remaining runs in the $\alpha^*$ column.

*Core violation probability:* We also report the percentage of runs in which a core violation exists (i.e., there exists at least one subset of authors who can all strictly improve by deviating from the current assignment). We refer to this as the *core violation probability* (CV-Pr).

*Social welfare:* Finally, we also measure the utilitarian and egalitarian social welfare (USW and ESW) defined above, which are the objectives maximized by TPMS and PR4A, respectively.

**Results.** The results are presented in Table 10.1. As expected, TPMS and PR4A achieve the highest USW and ESW, respectively, on all datasets because they are designed to optimize these objectives. In CVPR 2017, CoBRA and TPMS always end up with zero ESW because this dataset includes many zero similarity scores, but PR4A is able to achieve positive ESW. In all datasets, CoBRA achieves a relatively good approximation with respect to USW, but this is not always the case with respect to ESW.

For example, in CVPR 2018, CoBRA achieves $0.004$ ESW on average whereas PR4A achieves $0.099$ ESW on average. This may be due to the fact that this dataset also contains many zero similarity scores, and the myopic process of CoBRA locks itself into a bad assignment, which the global optimization performed by PR4A avoids.

While CoBRA suffers some loss in welfare, TPMS and PR4A also generate significant adverse incentives. They incentivize at least one community to deviate in almost every run of each dataset (CV-Pr). While the magnitude of this violation is relatively small when it is finite (except for TPMS in CVPR 2017), TPMS and PR4A also suffer from unbounded core violations in more than half of the runs for CVPR 2017; this may again be due to the fact that many zero similarity scores lead to deviations by groups where each agent has zero utility under the assignments produced by these algorithms.

Of all these results, the high probability of core violation under TPMS and PR4A is perhaps the most shocking result; when communities regularly face adverse incentives, occasional deviations may happen, which can endanger the stability of the conference. That said, CoBRA resolves this issue at a significant loss in fairness (measured by ESW). This points to the need for finding a middle ground where adverse incentives can be minimized without significant loss in fairness or welfare.

## 10.6 Discussion

The primary contribution of our work — introducing the core as a group fairness notion in the peer review setting — is conceptual in nature. This work introduces a novel notion of group fairness, called core, in a peer review setting which asks that each group is treated in a way that does not have any incentive to deviate and make its own conference. We show that a review assignment in the core always exists when each submission is authored by one agent, each agent serves as a reviewer and no author submits a number of submissions that require more reviewers to be reviewed than the number of submissions that the author can review. While, in the worst case, our algorithm achieves bad approximations with respect to often desirable desideratums, using real data, we show that in the average case, the approximations are quite good. On the other hand, famous reviewing assignment systems fail to satisfy group fairness very often and therefore incentivize communities to deviate from the current peer review process.

Our theoretical results serve merely as the first step toward using it to find reviewer assignments that treat communities fairly and prevent them from deviating. As such, our algorithm has significant limitations that must be countered before it is ready for deployment in practice. A key limitation is that it only works for single-author submissions, which may be somewhat more realistic for peer review of grant proposals, but unrealistic for computer science conferences. We also assume that each author serves as a potential reviewer; while many conferences require this nowadays, exceptions must be allowed in special circumstances. We also limit the number of submissions by any author to be at most $\lfloor k_a/k_p \rfloor$, which is a rather small value in practice, and some authors ought to submit more papers than this. We need to make this assumption to theoretically guarantee that a valid assignment exists. An interesting direction is to design an algorithm that can produce a valid assignment in the (approximate) core whenever it exists. Another limitation of our algorithm is that it assumes that each paper is assigned to the same number of reviewers and each reviewer can review at most the same number of papers. The extension of our algorithm for the multi-author case, and for the cases that papers are assigned to different numbers of reviewers and reviewers review different number of papers is an important open problem. Moreover, even if a community has incentives to deviate, the cost of deviating, may outweigh

the benefits of doing it. A more general model that also includes the cost for a community to break away from a large conference is another interesting direction. Finally, deploying group fairness in real-world peer review processes may require designing algorithms that satisfy it approximately at minimal loss in welfare, as indicated by our experimental results.

# Chapter 11

# Fair Team Formation

## 11.1 Introduction

The computer science department at University X is organizing a visit day for its newly admitted students. One of the most anticipated activity is the campus tour, during which the admitted students get to see the department they might one day join. Due to COVID-19 related capacity restrictions, the admitted students are divided into $k$ separate tours. But more tours means more volunteers. Luckily, $n$ current graduate students have agreed to help. We want to partition them almost equally between the $k$ tours so that all the admitted students have equal opportunity to socialize with the current students. However, the current students have developed a social network and, to ensure a pleasant experience to encourage them to volunteer in the future, we would like to assign volunteers to tours with as many of their friends as possible.

In this chapter, we introduce and study a model that captures such real-life applications. Specifically, we consider the problem of partitioning $n$ agents into $k$ almost equal-sized groups (each with size either $\lfloor n/k \rfloor$ or $\lceil n/k \rceil$), when the agents are connected via an undirected social network indicating *friendships*. An agent's utility for being part of a group is the number of her friends who are in that group.

Formally, this model sits within the *hedonic games* formalism in cooperative game theory with nontransferable utilities [Aziz and Savani, 2016]. Two compelling axiomatic guarantees that have received significant attention in this literature are *the core* [Gillies, 1953], which informally requires that there be no *deviating coalition* of agents such that each agent in the coalition has strictly more utility for the coalition than for her group in the given partition, and *envy-freeness* [George and Marvin, 1958], which informally requires that no agent receive strictly more utility when swapping places with another agent in the given partition. However, this literature typically does not impose any restriction on the partition (including on the number of groups it has). This would make our problem trivial because the grand coalition — a single group containing all agents — would trivially satisfy both the core and envy-freeness requirements. To study the core and envy-freeness meaningfully, this literature allows agents to have negative utility for other agents. We are interested in the case where the utilities are non-negative, but we require there to be exactly $k$ groups and the groups to be of approximately equal sizes.[1]

[1] In section *Beyond Balancedness*, we briefly consider imposing only the former restriction, allowing $k$ arbitrarily-sized non-empty groups.

This problem can be viewed as a multi-dimensional generalization of the *stable roommates problem* [Irving, 1985], in which the goal is to partition $2n$ agents between $n$ rooms of capacity 2 each, and agents have preferences over who they wish to have as a roommate. The core becomes a notion of stability: if a pair of agents prefer each other to their assigned roommates, they may actually deviate and rent a room by themselves. But the core has also been studied in contexts where groups cannot really deviate, such as allocation of public resources [Aziz et al., 2017, Fain et al., 2018, Conitzer et al., 2019] and clustering [Chen et al., 2019, Micha and Shah, 2020]. In such cases, it serves as a notion of *group fairness*, generalizing *proportionality* [Steinhaus, 1948], because it posits that each set of $\lfloor n/k \rfloor$ or $\lceil n/k \rceil$ agents is entitled to be a group on its own and demands that the agents in the set be treated no worse than if they were their own group. Envy-freeness, on the other hand, serves as a notion of *individual fairness* and requires that no agent envy another agent for the group they belong to.

## Our Results

For the core, we study bicriteria approximations of the form $(\alpha, \beta)$-core, where a deviating coalition must improve the utility of each of its members by more than a multiplicative factor of $\alpha$ and an additive factor of $\beta$.[2]

We begin with the most well-studied case of $k = 2$. We show that the $(2, 0)$-core is non-empty, and a 2-partition in the $(3, 0)$-core can be found in polynomial time. For larger $k$, we note that a $k$-partition in the $(1, k)$-core always exists when $n < k^2 + k$ and provide a lower bound proving that this guarantee is asymptotically the best possible with respect to an additive approximation. We show that a finite multiplicative approximation of the core is possible in general, and when $n \geqslant k^2 + k$, a min $k$-cut (a $k$-partition minimizing the so-called "cut size") is in the $(2k - 1, 0)$-core. While finding a min $k$-cut is known to be NP-hard, we present a polynomial-time algorithm that finds a (different) partition with the same approximation guarantee. We also show that min $k$-cuts cannot provide an asymptotically better approximation guarantee (i.e., our analysis is almost tight), and conjecture that no algorithm can.

For envy-freeness, we consider a similar additive approximation, EF-$r$, where an agent's utility must not increase by more than $r$ when swapping places with another agent. We make a connection to discrepancy theory [Chen et al., 2014] to show that a EF-$O(\sqrt{\frac{n}{k} \cdot \log k})$ partition always exists, and it can be computed efficiently. We conjecture that a EF-2 partition may always exists for any $k$.

Finally, in section *Beyond Balancedness*, we consider a classical variation of our model, where the only requirement is to create $k$ non-empty groups but the groups can be of arbitrary sizes. We study both the core and envy-freeness in this case and provide several tight approximation bounds.

## Related Work

Our work can be viewed as a hedonic game with symmetric, binary, and additively separable preferences, and with the restriction that the partition produced have exactly $k$ almost equal-sized parts. Hedonic games with restrictions on the number of groups have been studied before, but under other objectives, such as swap stability [Bilò et al., 2022], Pareto optimality [Cseh et al., 2019]. Moreover, Sless et al. [2018] consider a strictly weaker version of the core, without the size constraint, for which they provide positive results that do not carry over to our imbalanced case. As noted in the introduction, our model

[2]That is, an agent with utility $u$ must receive utility more than $\alpha u + \beta$ after deviating.

is a generalization of the stable roommates problem of partitioning $2n$ agents into $n$ pairs, where the widely studied notion of stability coincides with the core. In this problem, with asymmetric preferences a solution in the core does not always exist — unlike in the bipartite version, referred to as the *stable marriage problem*, in which it is guaranteed to exist [Gale and Shapley, 1962] — but can be found in polynomial time when it does [Irving, 1985]. When preferences are symmetric, however, a solution in the core always exists and can be found efficiently; for instance, one can repeatedly match and remove a pair of agents with the highest utility. The three-dimensional version of this problem — partitioning $3n$ agents into groups of size $3$ each — has also received significant attention. In this case, even with symmetric additive preferences, a solution in the core may not exist [Arkin et al., 2009], and checking whether it does is NP-hard [Chen and Roy, 2021]. However, if we further restrict the preferences to be binary, then McKay and Manlove [2021] show that a solution in the core always exists and can be found efficiently. Our problem can be seen as a multidimensional generalization of the roommate problem with symmetric binary additive preferences.

Envy-freeness has been studied recently in the hedonic games literature [Barrot and Yokoo, 2019], again with possibly negative utilities. Another concept similar to envy-freeness is *Nash-stability* [Bogomolnaia and Jackson, 2002, Olsen et al., 2012], which requires that no agent be happier by *joining* another part (rather than by swapping places with an agent in another part).[3] In our graph theoretic framework, this is equivalent to asking that each node have at least as many neighbors in its own part as in any other part. This has been studied extensively in graph theory using terms such as satisfactory partitions [Bazgan et al., 2010], friendly partitions [Aharoni et al., 1990], and internal partitions [Ban and Linial, 2016], but under only the restriction that each part is non-empty. This problem is also studied in the case where the parts are required to be of almost the same size [Bazgan et al., 2010]. However, since such partitions do not always exist, this literature primarily focuses on the computational complexity of checking the existence of such partitions and approximating the most satisfactory partitions.

Instead, our focus is on providing worst-case guarantees on the necessary violation of envy-freeness, as is commonly done in the literature on fair resource allocation [Lipton et al., 2004, Caragiannis et al., 2019b, Aziz et al., 2019b]. We make a connection to discrepancy theory [Chen et al., 2014] to establish an $O(\sqrt{n})$ bound. In discrepancy theory, the goal is to distribute each agent's friends as evenly as possible between the parts, so that not only does an agent not have many more friends in another part than her own part, she also does not have many more friends in her own part than in any other part. The latter restriction, a flipped version of the satisfactory partition problem, has also been studied separately as the co-satisfactory or unfriendly partition problem [Aharoni et al., 1990]. Manurangsi and Suksompong [2021] use discrepancy theory in a similar problem with $n$ agents partitioned into $k$ groups, but with the agents having utilities over goods being allocated to the groups, not over the other agents.

## 11.2 Preliminaries

For $t \in \mathbb{N}$, let $[t] = \{0, \ldots, t-1\}$. We consider a set $V = [n]$ of agents who are members of a social network. The network is represented by an undirected graph $G(V, E)$, where the agents are the nodes and an edge $(i, i') \in E$ indicates friendship between agents $i$ and $i'$. This induces the utility function of agent $i$, denoted $u_i : V \to \{0, 1\}$, where $u_i(i') = 1$ if $(i, i') \in E$ and $0$ otherwise. Let $N_G(i)$ denote the

[3]The two differ only when the other part consists entirely of the agent's friends.

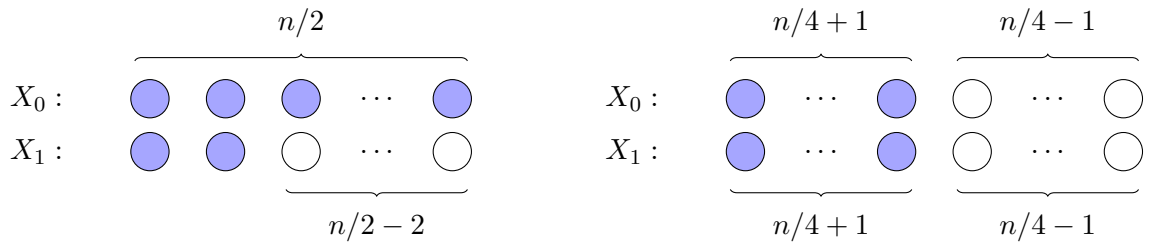


(a) 2-partition in the core

(b) EF1 2-partition

Figure 11.1: An instance with $K_{n/2+2}$ (blue nodes) and $n/2-2$ isolated agents (white nodes), where partitions that achieve the best approximation with respect to core achieve the worst approximation with respect to envy-freeness, and vice versa.

set of neighbors of agent $i$ in $G$, i.e., $N_G(i) = \{i' \in V : (i, i') \in E\}$. We refer to $d_G(i) = |N_G(i)|$ as the degree of agent $i$. We omit $G$ when it is clear from the context.

A *$k$-partition* of $V$ is given by $X = (X_0, \ldots, X_{k-1})$, where $X_j \cap X_{j'} = \emptyset$ for all distinct $j, j' \in [k]$; $X_j \neq \emptyset$ for all $j \in [k]$; and $\cup_{j\in[k]} X_j = V$. We may refer to an individual group $X_j$ as a *part*. With slight abuse of notation, we denote by $X(i)$ the part $X_j$ to which agent $i$ belongs (i.e., $i \in X_j$). We assume that $n \geqslant k$, so a $k$-partition exists. A $k$-partition is called *balanced* if $\lfloor n/k \rfloor \leqslant |X_j| \leqslant \lceil n/k \rceil$ for all $j \in [k]$, and called *imbalanced* otherwise. Hereinafter, whenever we refer to a $k$-partition, we mean a balanced $k$-partition, unless explicitly specified otherwise. The utility of agent $i$ for $S \subseteq V$ is denoted by, with slight abuse of notation, $u_i(S)$. We assume that utilities are additive, i.e., $u_i(S) = \sum_{i'\in S} u_i(i') = |S \cap N(i)|$.

In this work, we focus on two fairness criteria. The first one is the *core* which, informally, requires that there be no group of agents (coalition) of size $\lfloor n/k \rfloor \leqslant |S| \leqslant \lceil n/k \rceil$ such that every agent in the coalition prefers to be in that coalition than in her own part; such a coalition is called "blocking".

**Definition 24.** *Fix $\alpha \geqslant 1$ and $\beta \geqslant 0$. A coalition $S \subseteq V$ is called $(\alpha, \beta)$-blocking for a $k$-partition $X$ if*

$$u_i(S) > \alpha \cdot u_i(X(i)) + \beta$$

*for every $i \in S$. A $k$-partition $X$ is said to be in the* $(\alpha, \beta)$-core *if there is no $(\alpha, \beta)$-blocking coalition $S$ with $\lfloor n/k \rfloor \leqslant |S| \leqslant \lceil n/k \rceil$. When $\alpha = 1$ and $\beta = 0$, we simply use the terms blocking coalition, and core.*

Another fairness criterion we focus on is envy-freeness.

**Definition 25.** *For $r \geqslant 0$, a $k$-partition $X$ is called* envy-free up to $r$*, denoted EF$r$ or EF-$r$, if, for every pair of agents $i, i' \in V$, $u_i(X(i)) \geqslant u_i(X(i') \cup \{i\} \setminus \{i'\}) - r$. When $r = 0$, we simply refer to this as envy-freeness (EF).*

For the proof techniques we plan to use, we need the following additional terminology. The *cut size* of a $k$-partition $X$, denoted $\mathrm{cut}(X)$, is the number of edges between its different parts, i.e., $\mathrm{cut}(X) = |\{(i, i') \in E : X(i) \neq X(i')\}|$. A $k$-partition with the smallest cut size is called a *min $k$-cut*. Note that

$$\mathrm{cut}(X) = \sum_{i\in V} (|N(i)| - u_i(X(i))) = 2|\mathsf{E}| - \sum_{i\in V} u_i(X(i)).$$

Hence, min $k$-cut also maximizes the social welfare among all $k$-partitions. Some of our results show that such solutions also satisfy good approximations of the core. Give disjoint sets of nodes $A$ and $B$, $\mathsf{E}(A, B)$ denotes the set of edges with one endpoint in $A$ and the other in $B$.

Let us also introduce standard graph theory terminology. We denote by $K_n$, $K_{n_1,n_2}$ and $K_{n_1,n_2,n_3}$ the complete undirected graph of $n$ vertices; the complete bipartite graph with $n_1$ and $n_2$ vertices on the two sides; and the complete tripartite graph with $n_1$, $n_2$, and $n_3$ vertices on the three sides, respectively.

### Core vs. Envy-Freeness

While both core and envy-freeness are notions of fairness, they are quite different desiderata and it is easy to see that they are incompatible with each other in general. In fact, in our setting, there are instances in which every partition in the core achieves the worst possible approximation with respect to envy-freeness, and every partition that achieves the best approximations with respect to envy-freeness, achieves the worst possible approximation with respect to the core.

To see that, consider a graph that consists of the clique $K_{\frac{n}{2}+2}$ and $\frac{n}{2}-2$ isolated nodes, where $n$ is a multiple of 4 and $k=2$. The only 2-partitions in the core are the ones in which $\frac{n}{2}$ nodes from the clique form one group, say $X_0$, and the remaining nodes form the other group, say $X_1$. On the other hand, the 2-partitions that achieve the best approximation with respect to envy-freeness (in particular, EF1) are the ones in which each group contains $\frac{n}{4}+1$ nodes of the clique along with $\frac{n}{4}-1$ isolated nodes. These two types of partitions are illustrated in Figure 11.1.

Intuitively, the two partitions are complete opposites of each other: the former divides the clique as unequally as possible between the two groups, while the latter divides the clique exactly equally.The decision of which of the two notions (and correspondingly, partitions) is more desirable depends on the application at hand. Hence, we study both these notions separately in this chapter.

## 11.3 Core

In this section, we study $k$-partitions in the (approximate) core. We start from the important case that $k=2$, which has received particular attention in the literature on satisfactory partitions [Bazgan et al., 2010]. First, we point out an interesting open question:

**Open Question 1.** *Does every graph admit a 2-partition in the core?*

While the existence of 2-partitions in the core is unsettled, we show that the $(2,0)$-core is always non-empty, and in particular, contains every min 2-cut.

**Theorem 59.** *For $k=2$, a min 2-cut is in the $(2,0)$-core.*

*Proof.* Let $X=(X_0,X_1)$ be a min 2-cut. Suppose for contradiction that there exists a $(2,0)$-blocking coalition $S$ of size $\lceil n/2 \rceil$ or $\lfloor n/2 \rfloor$. Let $X_0^* = X_0 \cap S$ and $X_1^* = X_1 \cap S$.

For each agent $i \in X_0^*$, $i \in S$ implies $u_i(S) > 2 \cdot u_i(X_0)$, which in turn implies $|N(i) \cap X_1^*| > 2 \cdot |N(i) \cap X_0 \setminus X_0^*|$. Summing over all $i \in X_0^*$, we obtain

$$\mathsf{E}(X_0^*, X_1^*) > 2 \cdot \mathsf{E}(X_0^*, X_0 \setminus X_0^*).$$

Similarly, for each agent $i \in X_1^*$, we have $|N(i) \cap X_0^*| > 2 \cdot |N(i) \cap X_1 \setminus X_1^*|$. Summing over all $i \in X_1^*$, we get

$$\mathsf{E}(X_0^*, X_1^*) > 2 \cdot \mathsf{E}(X_1^*, X_1 \setminus X_1^*).$$

Combining the two equations, we have

$$\begin{aligned}\mathsf{E}(X_0^*, X_1^*) &> 2 \cdot \max\{\mathsf{E}(X_0^*, X_0 \setminus X_0^*), \\ &\qquad\qquad \mathsf{E}(X_1^*, X_1 \setminus X_1^*)\} \\ &\geqslant \mathsf{E}(X_0^*, X_0 \setminus X_0^*) + \mathsf{E}(X_1^*, X_1 \setminus X_1^*). \end{aligned} \tag{11.1}$$

Now, consider the 2-partition $X' = (S, V \setminus S)$. We will show that $\mathrm{cut}(X) > \mathrm{cut}(X')$, which will contradict $X$ being a min 2-cut. We have

$$\begin{aligned}\mathrm{cut}(X) &= \mathsf{E}(X_0, X_1) \\ &= \mathsf{E}(X_0^*, X_1^*) + \mathsf{E}(X_0^*, X_1 \setminus X_1^*) \\ &\quad + \mathsf{E}(X_1^*, X_0 \setminus X_0^*) + \mathsf{E}(X_0 \setminus X_0^*, X_1 \setminus X_1^*) \\ &\geqslant \mathsf{E}(X_0^*, X_1^*) + \mathsf{E}(X_0^*, X_1 \setminus X_1^*) \\ &\quad + \mathsf{E}(X_1^*, X_0 \setminus X_0^*) \\ &> \mathsf{E}(X_0^*, X_0 \setminus X_0^*) + \mathsf{E}(X_1^*, X_1 \setminus X_1^*) \\ &\quad + \mathsf{E}(X_0^*, X_1 \setminus X_1^*) + \mathsf{E}(X_1^*, X_0 \setminus X_0^*) \\ &= \mathrm{cut}(X'),\end{aligned}$$

where the strict inequality uses Equation (11.1). This is the desired contradiction. □

While Theorem 59 is a strong existential result, it does not come with an efficient algorithm as finding a min 2-cut (also known as the minimum bisection problem) is NP-hard [Garey and Johnson, 1979]. This leads to our next open problem:

**Open Question 2.** *Can a 2-partition in the $(2, 0)$-core be computed in polynomial time?*

As a slight remedy, we will later show that at least a 2-partition in the $(3, 0)$-core can be computed efficiently.

Next, we focus on $k \geqslant 3$ and show that the core can be empty in this case.

**Theorem 60.** *When $k \geqslant 3$, there exists an instance in which no $k$-partition is in the $(\alpha, 0)$-core for any $\alpha \geqslant 1$, and also there exists an instance in which no $k$-partition is in the $(1, \beta)$-core for $\beta < k/2 - 2$.*

*Proof.* Fix $k \geqslant 3$. For the first claim, consider a cycle with $n = k + 1 \geqslant 4$ nodes. Fix an arbitrary $k$-partition $X$. Note that $X$ must consist of one part with two nodes and $k - 1$ parts with a single node each. Without loss of generality, let $X_0$ be the part with $|X_0| = 2$. Note that in a cycle of length at least 4, the size of the smallest maximal matching is at least 2. Hence, there must exist agents $i, i' \notin X_0$ that are connected by an edge. Since the coalition $\{i, i'\}$ is allowed to deviate, they can both go from receiving utility 0 to receiving utility 1, implying that $X$ is not in the $(\alpha, 0)$-core for any $\alpha > 0$.

For the second claim, consider the graph $G$ formed by $k + 1$ disjoint cliques of size $k - 1$ each, denoted $C_0, \ldots, C_k$. Hence, $n = k^2 - 1$. Let $X$ be any $k$-partition of $G$. First, we claim that there

exists $\ell^* \in [k+1]$ such that $|C_{\ell^*} \cap X_j| \leqslant (k+1)/2$ for all $j \in [k]$. If this is not true, then for every $\ell \in [k+1]$, there exists at least one $j_\ell \in [k]$ with $|C_\ell \cap X_{j_\ell}| > (k+1)/2$. Note that such $j_\ell$ must be unique. Further, because $|X_j| \leqslant \lceil n/k \rceil \leqslant k+1$ for all $j \in [k]$, we must have $j_\ell \neq j_{\ell'}$ for distinct $\ell, \ell' \in [k+1]$. However, this is not possible as there are $k+1$ cliques but only $k$ parts. Now, for every agent $i \in C_{\ell^*}$, we have $u_i(X(i)) \leqslant (k-1)/2$. On the other hand, $C_{\ell^*}$ is a feasible deviating coalition as $|C_{\ell^*}| = k-1 = \lfloor n/k \rfloor$. Further, for every $i \in C_{\ell^*}$, we have $u_i(C_{\ell^*}) = k-2 \geqslant u_i(X(i)) + (k-3)/2$. Hence, $C_{\ell^*}$ is a $(1, k/2-2)$-blocking coalition, as desired. □

While above we show that the core can be empty when $k \geqslant 3$ , these examples are somewhat unsatisfactory as they crucially rely on $n$ not being divisible by $k$, which leads to another interesting open question:

**Open Question 3.** *Does every graph with $n$ nodes admit a $k$-partition in the core, if $k$ divides $n$?*

When $n < k^2 + k$, note that any deviating coalition has size at most $\lceil n/k \rceil \leqslant k+1$, which means that no agent can improve her utility by an additive factor of more than $k$ when deviating. Hence, every $k$-partition is trivially in the $(1, k)$-core.

**Proposition 12.** *When $n < k^2 + k$, every $k$-partition is in the $(1, k)$-core.*

From Theorem 60, we know that one cannot obtain a purely multiplicative guarantee of the form $(\alpha, 0)$-core for any $\alpha \geqslant 1$ and cannot obtain an additive guarantee of the form $(1, \beta)$-core for any $\beta \leqslant k/2 - 2$. Thus, we conclude that the guarantee in Proposition 12 is asymptotically tight with respect to these two ways of approximation when $n < k^2 + k$.

Next, we consider the case of $n \geqslant k^2 + k$. Interestingly, while a purely multiplicative approximation of the core cannot be guaranteed in general, we show that this is possible when we know $n \geqslant k^2 + k$. Specifically, the next result shows that a $k$-partition in the $(2k-1, 0)$-core always exists and can be found in polynomial time, if $n \geqslant k^2 + k$. In this case, we in fact show that every min $k$-cut is in the $(2k-1, 0)$-core, but we can also use a local search procedure, presented as Algorithm 18, to obtain the same approximation guarantee efficiently. At a high level, the algorithm works as follows. It starts from an arbitrary $k$-partition and draws a directed edge from agent $i$ to agent $i'$, with $X(i) \neq X(i')$, if $i$ envies $i'$ by more than one. If there exists a directed cycle, say $(i_0, i_1, \ldots, i_{s-1}, i_0)$, it shifts the agents along the cycle, i.e., $i_0$ is moved to $X(i_1)$, $i_1$ is moved to $X(i_2)$ and so on, while $i_{s-1}$ is moved to $X(i_0)$. Then, it updates the directed edges and continues eliminating cycles in this fashion. When there are no cycles left, it searches for pairs of agents, $i$ and $i'$, with $X(i) \neq X(i')$, such that $i'$ has zero utility for her group and a positive utility for $i$'s group, $i$ envies $i'$, and the envy is by more than one if $i$ and $i'$ are not neighbors. If such a pair exists, it swaps $i$ and $i'$. Then, the algorithm updates the directed edges, search for cycles or such pairs, and eliminates them until until no cycles or such pairs are left. Throughout the procedure, the cut size strictly decreases, and we establish that the approximation guarantee holds at any local minimum.

**Theorem 61.** *When $n \geqslant k^2 + k$, every min $k$-cut is in the $(2k-1, 0)$-core, and Algorithm 18 returns a $k$-partition in the $(2k-1, 0)$-core in polynomial time.*

*Proof.* First, we show that Algorithm 18 terminates in polynomial time by arguing that $\mathrm{cut}(X)$ strictly decreases in every iteration of the while loop. If we find a cycle in Line 4, then during the cyclic shift

**Algorithm 18** Local Min-Cut

1: $X \leftarrow$ an arbitrary $k$-partition
2: **while** true **do**
3: Build a directed graph $G' = (V', \mathsf{E}')$ with $V' = V$ and $\mathsf{E}' = \{(i, i') : u_i(X(i')) > u_i(X(i)) + 1\}$
4: **if** there is a cycle $(i_0, i_1, \ldots, i_{s-1}, i_0)$ in $G'$ **then**
5: **for** $\ell \in [s]$ **do**
6: $X(i_\ell) \leftarrow X(i_\ell) \setminus \{i_\ell\}$
7: $X(i_{\{\ell+1 \mod s\}}) \leftarrow X(i_{\{\ell+1 \mod s\}}) \cup \{i_\ell\}$
8: **end for**
9: **else if** $\exists (i, i')$ s.t. $u_{i'}(X(i')) = 0$ and $u_i(X(i')) > u_i(X(i))$ **then**
10: **if** $(i, i') \notin \mathsf{E}$ or $u_i(X(i')) > u_i(X(i)) + 1$ **then**
11: $X(i) \leftarrow X(i) \cup \{i'\} \setminus \{i\}$
12: $X(i') \leftarrow X(i') \cup \{i\} \setminus \{i'\}$
13: **end if**
14: **else**
15: break
16: **end if**
17: **end while**
18: **return** $X$

of nodes along this cycle, each node gains at least 1 utility. Since the social welfare strictly increases, $\mathrm{cut}(X)$ strictly decreases. Similarly, if we find two agents $i$ and $i'$ such that $i'$ has no neighbors in $X(i')$ but $i$ has at least two more neighbors in $X(i')$ than in $X(i)$, then swapping $i$ and $i'$ also strictly decreases the cut size. Further, if $i'$ is not a neighbor of $i$, then we only need $i$ to have at least one more neighbor in $X(i')$ than in $X(i)$. Hence, in any case, $\mathrm{cut}(X)$ strictly reduces in every iteration of the while loop, resulting in termination in polynomial time.

Let $X$ be either a min $k$-cut or the output of Algorithm 18. Suppose for contradiction that there is a $(2k-1, 0)$-blocking coalition $S$ of size $\lceil n/k \rceil$ or $\lfloor n/k \rfloor$. We first show the following lemma.

**Lemma 35.** *For $i \in S$, if $u_i(S \cap X_j) \leqslant u_i(X(i)) + 1$ for each $j \in [k]$, then $u_i(X(i)) = 0$.*

*Proof.* Suppose there exists $i \in S$ with $u_i(S \cap X_j) \leqslant u_i(X(i)) + 1$ for each $j \in [k]$ but $u_i(X(i)) \geqslant 1$. Then,

$$\begin{aligned} u_i(S) &= \textstyle\sum_{j \in [k]} u_i(S \cap X_j) \\ &\leqslant (k-1)(u_i(X(i)) + 1) + u_i(X(i)) \\ &\leqslant 2(k-1) \cdot u_i(X(i)) + u_i(X(i)) \\ &= (2k-1) \cdot u_i(X(i)), \end{aligned}$$

which contradicts $S$ being a $(2k-1, 0)$-blocking coalition. □

Suppose that there exists $i \in S$ such that $u_i(S \cap X_j) > u_i(X(i)) + 1$ for some $j \in [k]$. Let $G'$ be the directed graph constructed from $X$ according to Line 3 of Algorithm 18. Then, there must be an edge from $i$ to every node in $X_j$ in $G'$, as $u_i(X(i)) + 1 < u_i(S \cap X_j) \leqslant u_i(X_j)$. Further, since

$u_i(S \cap X_j) > 0$, $S \cap X_j \neq \emptyset$. Hence, $i$ has an edge to some node in $S$ in $G'$. Note that there can be no cycle in $G'$: if $X$ is the output of Algorithm 18, this would contradict the while loop terminating, and if $X$ is a min $k$-cut, a cyclic shift of nodes like in Algorithm 18 would reduce the cut size, which would be a contradiction. Since there is no cycle in $G'$, consider the longest path in $G'$ starting at $i$ and only containing nodes in $S$. Suppose it is $(i, i_1, \ldots, i_t, i')$. Then, $i'$ must satisfy the condition of Lemma 35, otherwise by the same reasoning as before, there would exist $j' \in [k]$ such that $S \cap X_{j'} \neq \emptyset$ and $i'$ has edges to all nodes in $X_{j'}$ in $G'$. This would either lead to a cycle or a longer path in $G'$ starting at $i$ and only containing nodes in $S$, which is a contradiction. Since $i'$ satisfies the condition of Lemma 35, we have $u_{i'}(X(i')) = 0$. We also have $u_{i_t}(X(i')) > u_{i_t}(X(i_t)) + 1$. If $X$ is returned by Algorithm 18, we get a contradiction because Algorithm 18 would have continued by swapping $i_t$ and $i'$ in Line 9. If $X$ is a min $k$-cut, then swapping $i_t$ and $i'$ would reduce the cut size, which would be a contradiction.

We have established that all $i \in S$ satisfy the condition from Lemma 35. Hence, $u_i(X(i)) = 0$ for all $i \in S$. But, since $n \geqslant k^2 + k$, we have $|S| \geqslant \lfloor n/k \rfloor \geqslant k + 1$, which implies that there exist $i_1, i_2 \in S$ with $X(i_1) = X(i_2)$. If $(i_1, i_2) \in \mathsf{E}$, then we have a contradiction with $u_{i_1}(X(i_1)) = u_{i_2}(X(i_2)) = 0$. Otherwise, there exists $i_3 \in S$, with $X(i_3) \neq X(i_1)$ and $(i_1, i_3) \in \mathsf{E}$. But, then we would have that either $(i_2, i_3) \notin E$ or $u_{i_3}(X_2) > u_{i_3}(X_3) + 1$, which means $i_2$ and $i_3$ would have been swapped under Algorithm 1, and hence this partition would not be a min cut. Thus, there is no such $(2k-1, 0)$-blocking coalition $S$.

□

In the proof of Lemma 35, note that if we assumed the deviating coalition $S$ to be a $(k, k-1)$-blocking coalition, then we would obtain a contradiction regardless of whether $u_i(X(i)) = 0$ or $u_i(X(i)) \geqslant 1$. Since the next part of the proof, which establishes that all $i \in S$ must satisfy the condition of Lemma 35, does not assume $n \geqslant k^2 + k$, we have that Algorithm 18 always finds a solution in the $(k, k-1)$-core. In particular, for $k = 2$, we can efficiently guarantee $(3, 1)$-core. Recall that Theorem 59 provides a slightly better guarantee of $(2, 0)$-core, but not in polynomial time.

**Corollary 62.** *For $k = 2$, Algorithm 18 returns a 2-partition in the $(3, 1)$-core in polynomial time.*

While it is an open question whether Algorithm 18 provides the best possible guarantee, we show, using an intricate instance, that our approximation analysis of min $k$-cuts in Theorem 61 is almost tight from a multiplicative point of view. Missing proofs can be found in the full version.

**Theorem 63.** *For $k \geqslant 3$, there exists an instance with $n \geqslant k^2 + k$ in which some min $k$-cut is not in the $(\alpha, 0)$-core for $\alpha < 2k - 2$.*

*Proof.* Consider a graph $G$ which consists of two disjoint cliques of size $2k - 3$ each, denoted by $A_1$ and $A_2$; $k - 1$ further disjoint cliques of size $2k - 4$ each, denote by $B_1, \ldots, B_{k-1}$; and another disjoint clique of size 2, denoted by $C$. Note that $n = 2(2k-3) + (k-1)(2k-4) + 2 = 2k^2 - 2k \geqslant k^2 + k$ for $k \geqslant 3$. We start with the following lemma.

**Lemma 36.** *If $X$ is a balanced min $k$-cut of $G$ such that for some $j^* \in [k]$, there exist $V \subseteq X_{j^*}$ and a clique $V' \subseteq \cup_{j \in [k]} X_j \setminus X_{j^*}$ with $\mathsf{E}(V', X_j \setminus V') = 0$ for every $j \in [k] \setminus j^*$, $\mathsf{E}(V', V) = 0$, $\mathsf{E}(X_{j^*} \setminus V, V') \geqslant \mathsf{E}(X_{j^*} \setminus V, V)$, and $|V| = |V'|$, then swapping the nodes between $V$ and $V'$, using an arbitrary bijection, does not increase the cut size.*

*Proof.* Since $V'$ is a clique, we easily see that the edges with two endpoints in different parts, except for part $X_{j^*}$ are not increased. On the other hand, as $\mathsf{E}(V', X_j \setminus V') = 0$ for any $j \in [k] \setminus j^*$, $\mathsf{E}(V', V) = 0$ and $\mathsf{E}(X_{j^*} \setminus V, V') \geqslant \mathsf{E}(X_{j^*} \setminus V, V)$, we see that all the edges with one endpoint to $X_{j^*}$ and the other endpoint to another part are not increased. Hence, cut is not increased. □

Let $X = (X_0, ..., X_{k-1})$ be an arbitrary balanced min $k$-cut of $G$. Suppose that the nodes of $A_1$ are spread among different parts. Then, there exists a part $X_{j_1}$ that contains at least two nodes of $A_1$, as $2k - 3 > k$ when $k > 3$. Let $S_1 = X_{j_1} \cap A_1$ and $V = X_{j_1} \setminus (S_1 \cup \{i_1\})$ where $i_1$ is an arbitrary node in $X_{j_1} \setminus S_1$ (such a node always exists as $n/k > |A_1| - 1 = 2k - 4$). Notice that $\overline{S}_1 = A_1 \setminus S_1$ is a clique such that $\mathsf{E}(\overline{S}_1, X_j \setminus \overline{S}_1) = 0$ for any $j \in [k] \setminus j_1$ and $\mathsf{E}(\overline{S}_1, V) = 0$. Moreover, notice that $|\overline{S}_1| = |V|$ as $|A_1| = 2k - 3$ and $|X_{j_1}| = 2k - 2$. We also see that

$$\mathsf{E}(S_1 \cup \{i_1\}, X_{j_1} \setminus (S_1 \cup \{i_1\})) =$$
$$\mathsf{E}(S_1, X_{j_1} \setminus (S_1 \cup \{i_1\})) + \mathsf{E}(i_1, X_{j_1} \setminus (S_1 \cup \{i_1\})) \leqslant 0 + |V|$$

while

$$\mathsf{E}(S_1 \cup \{i_1\}, \overline{S}_1) \geqslant 2 \cdot |\overline{S}_1|$$

as $|S_1| \geqslant 2$ and all the agents in $S_1$ are connected with all the agents in $\overline{S}_1$. From Lemma 36, we get that if we swap the nodes among $V$ and $\overline{S}_1$, using an arbitrary mapping, the cut of the partition is not increased. Hence, there exists a balanced min $k$-cut in which all the nodes in $A_1$ are assigned to the same part $X_{j_1}$. Given this partition, suppose that the nodes of $A_2$ are spread among different parts. Then, there exists a part $X_{j_2}$ different than $X_{j_1}$ that contains at least two nodes of $A_2$, and by a similar argument as above we conclude in a balanced min $k$-cut in which all the nodes in $A_1$ are assigned to $X_{j_1}$ and all the nodes in $A_2$ are assigned to $X_{j_2}$. Now, starting from this partition, suppose that the nodes of $B_1$ are spread among different parts. Then, as $|B_1| = 2k - 4$, there exists a part $X_{j'_1}$ different that $X_{j_1}$ and $X_{j_2}$ which contains at least two nodes from $B_1$, when $k > 2$. Hence, by the same arguments as above we can conclude in a balanced min $k$-cut, in which all the nodes in $A_1$ are assigned to $X_{j_1}$, all the nodes in $A_2$ are assigned to $X_{j_2}$, and all the nodes in $B_1$ are assigned to $X_{j'_1}$. By continuing this way and as $|B_\ell| = 2k - 4$ for any $\ell \in \{1, ..., k-1\}$, we conclude that there exists a balanced min $k$-cut, in which all the nodes in $A_1$ are assigned to $X_{j_1}$, all the nodes in $A_2$ are assigned to $X_{j_2}$, and all the nodes in $B_\ell$, for $\ell \in \{1, ..., k-3\}$, are assigned to $X_{j'_\ell}$. Now, the last part different that $X_{j_1}$, $X_{j_2}$ and $X_{j'_\ell}$ for $\ell \in \{1, ..., k-3\}$, denoted by $X_{j'_{k-2}}$, contains at least two nodes from $B_{k-2}$ or $B_{k-1}$. Without loss of generality, we assume that $|X_{j'_{k-2}} \cap B_{k-2}| \geqslant 2$. By doing the same arguments as before, we find a balanced min $k$-cut partition $X' = (X'_0, ..., X'_{k-1})$ in which all the nodes of each clique $A_1$, $A_2$, $B_1$,...,$B_{k-2}$ are assigned to the same group. Let $Q = A_1 \cup A_2 \cup B_1 \cup ... \cup B_{k-2}$. Then notice that for each $j \in [k]$, $|X'_j \cap Q| \geqslant 2k - 4$. Hence, the nodes of $B_{k-1}$ and $C$ are spread across different parts and each part contains at most two nodes of $B_{k-2}$.

Now, we can see that if $C = \{c_1, c_2\}$ and $B_{k-1} = \{b_1, \ldots, b_{2k-4}\}$, then $X''$ given by $X''_0 = A_1 \cup \{c_1\}$, $X''_1 = A_2 \cup \{c_2\}$, and $X''_{j+1} = B_j \cup \{b_{2j-1}, b_{2j}\}$ for $j \in \{1, \ldots, k-2\}$ is a balanced min $k$-cut. But then, the coalition $S = C \cup B_{k-1}$ is a $(2k-3, 0)$-blocking coalition, as the utility of each agent in $C$ increases by an infinite multiplicative factor when deviating with $S$, while that of each agent in $B_{k-1}$ increases by a multiplicative factor of $2k - 3$ when deviating with $S$. □

## 11.4 Envy-Freeness

We now turn our attention to finding $k$-partitions that are (approximate) envy-free. We start by showing that EF1 cannot always be guaranteed.

**Theorem 64.** *Even when $k = 2$, a 2-partition that is EF1 does not always exist.*

*Proof.* Consider the $K_{3,3,3}$ graph which consists of three set of three nodes each, denoted by $C_1 = \{c_{11}, c_{12}, c_{13}\}$, $C_2 = \{c_{21}, c_{22}, c_{23}\}$ and $C_3 = \{c_{31}, c_{32}, c_{33}\}$, respectively, and every node of each set is adjacent to every node in the other two sets.

For the sake of contradiction, assume that $X = (X_0, X_1)$ is a partition of the graph that is EF1. Since the graph is 6-regular, we can see that $|X_0| \geqslant 4$ and $|X_1| \geqslant 4$, as if an agent $i$ is assigned to a part with only at most two of its neighbours, then the other four of its neighbours are assigned to the other part along with an agent $i'$ which is not neighbour of $i$, and then $i$ envies $i'$ for more than one agent. Without loss of generality, we assume that $|X_0| = 4$. If $X_0$ contains three nodes of the same set, then we can easily see that this partition is not EF1, as each of them is assigned to the same group with at most one of its neighbours. As there are three sets and $X_0$ contains four agents, we see that two agents of the same set, say $c_{11}$ and $c_{12}$, are assigned to $X_0$. Then these two agents are in the same part along with at most two of its neighbours, while all the remaining nodes are assigned to $X_1$. Then, $c_{11}$ and $c_{12}$ envy $c_{13}$ for more than one agents, which is a contradiction. □

To obtain non-trivial approximations to envy-freeness for higher values of $k$, that too via partitions, we turn to the literature on discrepancy theory. Intuitively, we want to color the elements of a set using $k$ colors such that each pre-determined subset has an approximately equal number of elements of each color. Formally, we are given a universe $\Omega = [n]$ and a set system $\mathcal{S} = \{S_0, \ldots, S_{m-1}\}$, where $S_i \subseteq [n]$ for each $i \in [m]$. The $k$-color discrepancy of a coloring $\chi : \Omega \to [k]$ on the set system $\mathcal{S}$ is defined as

$$\mathrm{disc}_k(\mathcal{S}, \chi) = \max_{j \in [k], i \in [m]} \left| \left| \chi^{-1}(j) \cap S_i \right| - |S_i|/k \right|.$$

The $k$-discrepancy of $\mathcal{S}$ is then the minimum $k$-color discrepancy over all $\chi$ is given as

$$\mathrm{disc}_k(\mathcal{S}) = \min_{\chi : \Omega \to [k]} \mathrm{disc}_k(\mathcal{S}, \chi)$$

. A celebrated result from this literature is that $\mathrm{disc}_k(\mathcal{S}) = O(\sqrt{\frac{n}{k} \ln(km/n)})$ for any set system $\mathcal{S}$ and a $k$-coloring achieving this bound can be computed in polynomial time [Chen et al., 2014, Corollary 44].

In our setting, with $\Omega = V = [n]$, a $k$-coloring $\chi : \Omega \to [k]$ induces a $k$-partition $X$ given by $X_j = \chi^{-1}(j)$ for all $j \in [k]$.[4] Further, if we consider the set system $\mathcal{S}$ where $S_i = N_G(i)$ for each $i \in [n]$ (i.e., with $m = n$), then we are guaranteed that agent $i$ can have at most $2\mathrm{disc}_k(\mathcal{S}, \chi)$ more neighbors in any other part than in her own part, implying EF-$2\mathrm{disc}_k(\mathcal{S}, \chi)$. The above discrepancy bound then immediately yields the existence of a $k$-partition that is EF-$O(\sqrt{\frac{n}{k} \ln k})$. However, this may not be balanced.

To fix this, we add another set $S_n = V$ to our set system; we now have $m = n + 1$, which does not asymptotically change the discrepancy bound. Now, we obtain a $k$-partition $X$ that is also approximately

[4]Technically, we also need to ensure $X_j \neq \emptyset$, but this is guaranteed due to the discrepancy bound.

balanced: $||X_j| - |X_{j'}|| = O(\sqrt{\frac{n}{k}\ln k})$ for all $j, j' \in [k]$. By arbitrarily moving $O(\sqrt{\frac{n}{k}\ln k})$ agents between parts, we can make it perfectly balanced, while only increasing the EF approximation by $O(\sqrt{\frac{n}{k}\ln k})$. Thus, we get the following.

**Theorem 65.** *For any $k \geqslant 2$, a $k$-partition that is EF-$O(\sqrt{\frac{n}{k}\ln k})$ is guaranteed to exist and can be computed in polynomial time.*

For discrepancy, the aforementioned upper bound is known to be almost tight: there is a lower bound of $\Omega(\sqrt{n/k})$ [Chen et al., 2014, Theorem 61]. However, for our "one-sided" envy-freeness guarantee, achieving a constant approximation remains an open question.

**Open Question 4.** *Does every graph admit a $k$-partition that is EF2, for all $k \geqslant 2$?*

## 11.5 Beyond Balancedness

An interesting variation of our problem is to drop the requirement of balancedness and simply seek $k$ non-empty groups, i.e., imbalanced $k$-partitions. This variation was first introduced by Gerber and Kobler [2000] and, since then, it has been given much attention [Bazgan et al., 2010] due to its importance to real-life applications such as clustering [Flake et al., 2004, Shafique, 2004].

In this section, we briefly consider this case and study both the core and envy-freeness for imbalanced $k$-partitions. For the core, we provide matching upper and lower bounds. For envy-freeness, we provide a complete picture for $k = 2$ by making a connection to the literature on satisfactory partitions, and point out interesting open questions for $k \geqslant 3$.

### 11.5.1 Core

Recall that core requires that there be no group of agents (coalition) such that every agent in the coalition prefers to be in that coalition than in her own part. In general, there is no direct correlation between the size of a coalition and the ease with which it can be blocking.[5] Hence, in the imbalanced case, we impose the same restriction on the size of a deviating coalition as we have on the size of a part in an imbalanced $k$-partition. Note that all parts in an imbalanced $k$-partition $X$ are required to be non-empty, which implies $1 \leqslant |X_j| \leqslant n - k + 1$ for all $j \in [k]$; hence, we require a deviating coalition $S$ to also satisfy $1 \leqslant |S| \leqslant n - k + 1$. This gives rise to the following variant of the core.

**Definition 26.** *Fix $\alpha \geqslant 1$ and $\beta \geqslant 0$. A coalition $S \subseteq V$ is called $(\alpha, \beta)$-blocking for an imbalanced $k$-partition $X$ if*

$$u_i(S) > \alpha \cdot u_i(X(i)) + \beta$$

*for every $i \in S$. An imbalanced $k$-partition $X$ is said to be in the* $(\alpha, \beta)$-imbalanced core *if there is no $(\alpha, \beta)$-blocking coalition $S$ with $1 \leqslant |S| \leqslant n - k + 1$. When $\alpha = 1$ and $\beta = 0$, we simply use the terms blocking coalition, and imbalanced core.*

[5] Smaller coalitions have the advantage of only having to improve the utility of fewer agents, whereas larger coalitions can include more friends of their members.

Note that the differing size restrictions on the deviating coalitions technically makes our results for the core under imbalanced $k$-partitions incomparable to our results for the core under balanced $k$-partitions.

We show that for any $k$, an imbalanced partition in the $(1, k-2)$-imbalanced core always exists and can be found in polynomial time.

**Algorithm 19** (Approximate) Imbalanced Core

**if** $G$ is connected **then**
  **for** $r = 0, \ldots, k-2$ **do**
    $i_r \leftarrow$ a leaf node in a spanning tree of $G \setminus \cup_{t\in[r]} X_t$
    $X_r \leftarrow \{i_r\}$
  **end for**
**else**
  $X_0 \leftarrow$ any connected component of $G$
  **for** $r = 1, \ldots, k-2$ **do**
    **if** $V \setminus \cup_{t\in[r]} X_t \neq \emptyset$ **then**
      $X_r \leftarrow \{i\}$, for an arbitrary $i \in V \setminus \cup_{t\in[r]} X_t$
    **else**
      $X_r \leftarrow \{i\}$, for an arbitrary $i \in X_0$
      $X_0 \leftarrow X_0 \setminus \{i\}$
    **end if**
  **end for**
**end if**
**if** $V \setminus \cup_{t\in[k-1]} X_t \neq \emptyset$ **then**
  $X_{k-1} = V \setminus \cup_{t\in[k-2]} X_t$
**else**
  $X_{k-1} = \{i\}$, for an arbitrary $i \in X_0$
  $X_0 \leftarrow X_0 \setminus \{i\}$
**end if**
**return** $X = (X_0, \ldots, X_{k-1})$

**Theorem 66.** *When $k \geqslant 2$, we can find an imbalanced $k$-partition in the $(1, k-2)$-imbalanced core in polynomial time, and in particular, when $k = 2$, we can efficiently find an imbalanced 2-partition in the imbalanced core. Moreover, when $k \geqslant 3$, there exists an instance in which no imbalanced $k$-partition is in the $(1, \beta)$-imbalanced core for any $\beta < k-2$.*

*Proof.* We show Algorithm 19 finds a 2-partition in the imbalanced core in polynomial time, while when $k > 2$, the same algorithm finds a partition in the $(1, k-2)$-imbalanced core.

For contradiction, assume that there is a blocking coalition $S$ for the $k$-partition $X$ computed by Algorithm 19 with $1 \leqslant |S| \leqslant n-k+1$ each of whose agents increased their utility by at least an additive factor of $k-1$.

First, we suppose that $G$ is connected. Notice that every connected graph admits a spanning tree, and that the graph stays connected when deleting a leaf from this tree. Hence, Algorithm 19 is well-defined in this case, and we obtain the guarantee that $X_{k-1}$ is a connected subgraph of $G$.

We claim that $S \cap X_{k-1} \neq \emptyset$ and $X_{k-1} \setminus S \neq \emptyset$. To see the former claim, note that if $|S| \leqslant k-1$, then $u_i(S) \leqslant k-2 \leqslant u_i(X(i)) + k - 2$ for any $i \in S$, which would be a contradiction. Hence, we must have $|S| \geqslant k$, which implies $S \cap X_{k-1} \neq \emptyset$. To see the latter claim, note that $|S| \leqslant n-k+1$. Also, $|X_{k-1}| = n-k+1$. Thus, if $S \supseteq X_{k-1}$, then we would have $S = X_{k-1}$. This would imply $u_i(S) = u_i(X(i))$ for all $i \in S$, which would again be a contradiction. Hence, we must have $X_{k-1} \setminus S \neq \emptyset$.

Fix $i_1^* \in S \cap X_{k-1}$ and $i_2^* \in X_{k-1} \setminus S$. Because $X_{k-1}$ is a connected subgraph of $G$, there exists a path from $i_1^*$ to $i_2^*$ using only the nodes in $X_{k-1}$. Consider the first edge of this path to travel out of $S$; say this edge is $(i', i'')$ with $i' \in S \cap X_{k-1}$ and $i'' \in X_{k-1} \setminus S$. When deviating from $X_{k-1}$ to $S$, agent $i'$ loses at least one neighbor (namely $i''$) from $X_{k-1}$ and may gain up to $k-1$ neighbors (the nodes in $\cup_{r \in [k-1]} X_r$). This implies $u_{i'}(S) \leqslant u_{i'}(X_{k-1}) + k - 2$, which is a contradiction.

Next, suppose $G$ is not connected. Since no connected component can contain all nodes of $G$, the algorithm must have moved at most $k-2$ nodes from $X_0$ to $\cup_{t \in \{1,\ldots,k-1\}} X_t$. Hence, none of the agents who are in $X_0$ in the final solution can join coalition $S$ as their utility cannot improve by more than an additive factor of $k-2$ when doing so. Further, if there exists $i \in S \cap X_{k-1}$, then $u_i(S) \leqslant u_i(X_{k-1}) + \sum_{r=1}^{k-2} u_i(X_r) \leqslant u_i(X_{k-1}) + k - 2$ (as we have already established $S \cap X_0 = \emptyset$), which is again a contradiction. Hence, we must have $S \subseteq \cup_{r \in \{1,\ldots,k-2\}} X_r$, implying that $|S| \leqslant k-2$. But then, $u_i(S) \leqslant k-3 < u_i(X(i)) + k - 2$ for all $i \in S$, which is again a contradiction.

For the lower bound, consider the complete graph $K_n$ with $n \geqslant k \cdot (k-1)$. Let $X$ be any $k$-partition of this graph. Due to the pigeonhole principle, there exists $r^* \in [k]$ such that $|X_{r^*}| \geqslant n/k \geqslant k-1$. Hence, the coalition $S = \cup_{r \in [k] \setminus \{r^*\}} X_j$ is allowed to deviate as $|S| \leqslant n-k+1$. Since each $X_r$ part of this coalition is non-empty, we have $u_i(S) \geqslant u_i(X(i)) + k - 2$ for each $i \in S$, implying that $X$ is not in the $(1, \beta)$-core for any $\beta < k-2$. □

Recall that in the proof of Theorem 60, we used an example with $n = k+1$ to establish that there may not exist any (balanced) $k$-partitions in the $(\alpha, 0)$-core for any $\alpha \geqslant 1$ or $(1, \beta)$-core for any $\beta < k/2 - 2$. Because all imbalanced $k$-partitions are also balanced for $n = k+1$, the imbalanced core becomes equivalent to the core. Hence, the negative results of Theorem 60 carry over to the imbalanced case, though the additive lower bound in Theorem 66 is better.

This shows that the guarantee in Theorem 66 is tight in two ways: one can hope for neither a purely multiplicative approximation of the form $(\alpha, 0)$-imbalanced core for any $\alpha \geqslant 1$, nor a better additive approximation of the form $(1, \beta)$-imbalanced core for any $\beta < k-2$.

### 11.5.2 Envy-Freeness

Finally, we turn our attention to envy-freeness. First, we use the following result from the literature on satisfactory partitions, restated in our framework, to establish the existence of an EF-2 imbalanced partition when $k = 2$.

**Theorem 67** (Stiebitz 1996, Bazgan et al. 2007)**.** *Given a graph $G = (V, E)$ and functions $a, b : V \to \mathbb{N}$ such that $d(i) \geqslant a(i) + b(i) + 1$ for every $i \in V$, there exists an imbalanced 2-partition $X = (X_0, X_1)$ of $V$ such that $u_i(X_0) \geqslant a(i)$ for each $i \in X_0$ and $u_i(X_1) \geqslant b(i)$ for all $i \in X_1$, and it can be computed in polynomial time.*

In our case, we use functions $a(i) = b(i) = \lfloor (d(i) - 1)/2 \rfloor$ for all $i \in V$. Note that these satisfy the condition $d(i) \geqslant a(i) + b(i) + 1$. Hence, the above result allows us to efficiently compute a 2-partition $X$ satisfying $u_i(X(i)) \geqslant \lfloor (d(i) - 1)/2 \rfloor$ for all $i \in V$. Since there are only two parts, this also implies that for all $i, i' \in V$,

$$\begin{aligned} u_i(X(i')) - u_i(X(i)) &\leqslant d(i) - 2 \cdot \lfloor (d(i) - 1)/2 \rfloor \\ &\leqslant d(i) - 2 \cdot (d(i) - 2)/2 = 2, \end{aligned}$$

which implies that $X$ is EF-2.

**Corollary 68.** *An imbalanced 2-partition that is EF-2 always exists and can be computed in polynomial time.*

Theorem 67 admits an extension to $k > 2$ parts, but in our case, this only guarantees that $u_i(X(i)) \geqslant \lfloor (d(i) - k + 1)/k \rfloor$ for all $i \in V$ [Bazgan et al., 2007]. This does not meaningfully limit the number of neighbors that agent $i$ has in another part and, therefore, fails to provide a non-trivial approximation to envy-freeness. That said, if one is interested in the slightly weaker guarantee of proportionality [Steinhaus, 1948], which, in our setting, would require $u_i(X(i)) \geqslant d(i)/k$, then this would provide an additive 1-approximation.

For the satisfactory partition problem, where the goal is to indeed minimize $u_i(X(i')) - u_i(X(i))$, as in the equation above, it is easy to see that an additive error of 2 is the best possible. Consider dividing any clique with an odd number of nodes into two parts. An agent $i$ in the smaller part will have at least two more neighbors in the larger part than in her own part. However, this does not hold for envy-freeness: if $i$ envisions *swapping places* with an agent $i'$ from the other part, then $X(i') \cup \{i\} \setminus \{i'\}$ will only contain one more neighbor of $i$ than $X(i)$ does. Nonetheless, notice that the example that is used in the proof of Theorem 64 can also be used to show that EF-1 cannot always be guaranteed even in the imbalanced case when $k = 2$.

## 11.6 Discussion

In this chapter, we considered the problem of partitioning $n$ agents into $k$ almost equal-sized groups, when the agents have binary preferences, induced by a social network. We designed algorithms which approximately satisfy two axiomatic fairness guarantees: the core and envy-freeness. Our work offers a number of exciting open questions. For example, is the core always non-empty when $k = 2$ or when $k$ divides $n$? Does an EF-2 partition always exist?

There are natural ways to extend our model. One can consider more general preferences than symmetric and binary. Symmetric weighted preferences are particularly interesting as while one can verify that our positive result of min 2-cut carries over this case, our guarantees of min $k$-cut for $k \geqslant 3$ are not easily expandable beyond the binary case. Moreover, if the agents are described by a number of attributes, the construction of fair and diverse groups is another interesting direction.

# Part IV

# Conclusion

In this thesis, we discussed how by building on ideas from the intersection of Economics and Computer Science, we can optimize decision-making processes that foster social good, analyze and propose enhancements to methods of civic participation within our existing democratic systems, and introduce innovative concepts of fairness and provably fair algorithms in various domains. Below, we summarize the different ideas and techniques from social choice theory and theoretical computer science that were used across different chapters. We also make connections between different chapters, by highlighting the interrelatedness of problems and emphasizing the application of the same ideas and techniques across different contexts. Finally, we discuss our future vision within these three categories—optimizing decision-making processes for social good, advancing civic participation, and pushing the boundaries of fairness using notions of fairness from economics.

# Ideas, Techniques and Connections Across Different Chapters

In Part I, we focused on two real life applications in which the goal is to make decisions that optimize the social good. In both chapters, we defined the social good as the (utilitarian) social welfare, a fundamental concept originating from social choice theory, which is used to evaluate the overall well-being within a society. In the first chapter of this part, we focused on the problem of allocating a limited budget of pooled diagnostic test during an epidemic. The first result indicates that by focusing on logistically simple solutions, we achieve a good approximation to the optimal solution. The second main result indicates that while the problem is in general computationally difficult, a simple greedy algorithm gives a good approximation. Therefore, while the problem falls into the category of resource allocation problems that have been traditionally studied in social choice theory, our results build upon the central idea of approximation algorithms, which is foundational in the field of theoretical computer science, and the utilization of greedy algorithms, which is another fundamental concept in theoretical computer science.

In Part II, we focused on democratic innovations for increasing civic participation. In the first chapter of this part, we considered sortition and designed algorithms for selecting panels that satisfy different properties, when the citizens lie in an underlying metric space. First, we asked for panels that minimize the social cost. Similarly to the concept of the social welfare, the notion of the (utilitarian) social cost stems from the realm of social choice and minimizing this objective aims to minimize the total burden on the society. Second, we asked for panels that are proportionally representative, where proportionality indicates that if there are $n$ citizens in total, and a panel should contain $k$ members, then proportionally every group of $n/k$ citizens is entitled to choose one panel member and, by extension, every group of $\ell \cdot n/k$ citizens is entitled to choose $\ell$ panel members for any $\ell$. We defined a panel to be proportionally representative if there is no group of citizens that has an incentive to deviate from the given panel and instead pick a panel with size equal to their entitlement which represents each of them more. This notion of proportional representation is inspired by the core, a central concept of stability in economics. Informally, the core is defined as the set of outcomes where no subgroup of individuals can deviate and form a more favorable outcome for itself. Here, the core is defined in accordance to the property of proportionality.

Under both desiderata, the main property of sortition that we asked to be preserved is that all the citizens have the same chance to be selected. This property differentiates sortition from a committee selection problem, where the goal is to select a representative committee without any further considera-

tion regrading the chances that a candidate has to be included in the committee. When the goal is to find panels that minimize the social cost, we show that in some cases by selecting a panel uniformly at random, we achieve a good approximation to the social cost, while in some other cases, we design selection algorithms that provide a trade-off between the optimal social cost and the minimum probability that each citizen is included in the panel. When the goal is to find panels that are in the core, we observed that there may not be any panel in the core, even in the absence of the constraint that all citizens be included with the same probability. Thus, we designed a greedy algorithm that achieves a constant approximation to the core. Again, in both cases we used approximation algorithms that are central in theoretical computer science, and in the second case we utilized greedy algorithms.

In the second chapter of Part II, we focused on a committee selection problem, where there is a set of $n$ voters and a set of candidates, each voter approves a subset of candidates, and the goal is to select a committee of size $k$. While most of the committee selection literature focuses on notions of group fairness, we proposed a novel notion of individual fairness and asked for committees that satisfy it. This notion of fairness again builds on the idea of proportionality, which suggests that at least $n/k$ voters are entitled to choose one committee member and, by extension, at least $\ell \cdot n/k$ voters are entitled to choose $\ell$ committee members for any $\ell$. We proved that there exist instances in which there are no committees that satisfy this property. Therefore, we turned our attention to the question of whether an approximation of this notion of fairness can be guaranteed. While we answered this question negatively for the general case, we designed a selection algorithm that achieves a constant approximation to this notion when the votes are structured. Hence, the idea of approximation algorithms is central to this chapter as well.

In the last chapter of Part II, we considered the question of how many sample votes are needed in the worst and average cases to learn the outcome of an election. While these kinds of questions are commonly studied in theoretical computer science, they are somewhat novel to social choice theory, which primarily focuses on studying different preference aggregation rules and properties of collective decision-making processes.

In Part III, we focused on algorithmic fairness and showed how notions of fairness that derive from economics and social choice theory can be very effective to domains beyond computational social choice including machine learning. In the first chapter of this part, we considered a clustering problem where there are $n$ data points and the goal is to select $k$ cluster centers from anywhere in the metric space, when each data point represents an agent whose cost for a clustering solution is equal to its distance from its closest cluster center. Here, we asked for proportionally fair solutions, where proportional fairness is defined in the same way that proportional representation was defined in the context of sortition. In particular, under proportionality every $n/k$ data points are entitled to choose one cluster center. Therefore, a cluster solution is proportionally fair if there is no subset of data points of size at least $n/k$ and a new potential cluster center such that each point in the subset is closer to the new center than to any center in the current solution. One might think that this is different from the core because we are not considering groups of size at least $\ell n/k$ (for some $\ell > 1$) deviating with $\ell$ new cluster centers. However, it is easy to see that this is actually equivalent to the definition of the core due to the definition of the cost function: if some $\ell n/k$ points were to deviate with $\ell$ new centers, at least one of the new centers would be the closest to at least $n/k$ of the deviating points, so these would also form a violation of proportional fairness. Thus, we in fact studied the core to the clustering problem in this chapter.

In Part II, we considered the sortition problem, where citizens lie in an underlying metric space, just as data points lie in a metric space in the clustering problem above. In both cases, $k$ representatives/cluster centers need to be chosen and each citizen/data point evaluates their distance to the

nearest representative/cluster center. Thus, the sortition problem can be viewed as the same as the clustering problem above, but with two additional restrictions: the cluster centers now must be data points themselves and each data point must be chosen as a cluster center with approximately equal probability. For both problems, we used modified versions of the greedy algorithm introduced by Chen et al. [2019]. In particular, for the clustering problem, we modified the algorithm to allow every point in a continuous metric space to be a possible cluster center (Chen et al. [2019] only considered taking as input a discrete set of possible cluster centers). For the sortition problem, we modified the same greedy algorithm but in a different way so that each citizen ("data point") is included in the panel ("clustering solution") with the same probability. Therefore, for both problems, the central idea of greedy algorithms plays a crucial role.

In the second chapter of Part III, we explored a multi-agent generalization of the classic multi-armed bandits problem, where pulling each arm generates different stochastic rewards to the different agents. The objective was to attain a policy, which is a distribution over arms, that ensures fairness among the agents. To accomplish this, we utilized the concept of maximizing Nash social welfare from social choice theory. To assess the performance of our proposed algorithms for this problem, we employed the notion of regret, a standard measure in theoretical computer science. Regret quantifies the deviation of an algorithm's performance from an optimal solution.

In the third chapter of Part III, we focused on two-sided markets, such as those present in recommender systems where a set of producers (of content or products) are matched/recommended to a set of consumers. We asked for solutions that are fair to both sides simultaneously. The notion of fairness used was envy-freeness, which derives from social choice theory and asks that no agent envy the allocation of another agent. In this context, the allocation of an agent is the set of agents from the other side that are matched to her. The idea of envy-freeness is central in fair division, and here we showed that it can be effective in a different domain (and applied to both sides of the market simultaneously).

In the fourth chapter of Part III, we delved into a peer review problem, where a set of agents review each others' submissions (e.g., as in a conference system), and each author submits a number of papers. The objective is to find an assignment of submissions to reviewers that satisfies the core. In this context, a review assignment is considered to be in the core if no subset of authors can assign a subset of their submissions among themselves and be happier. Unlike in the previous cases, there is no global resource (number of cluster centers or panel members) that is proportionally scaled with the group size. Instead, each author contributes an endowed resource (their reviewing capacity) and a group can deviate with their own contributed resource, which serves as their entitlement. In this sense, the core captures a sense of stability, and we introduced a greedy algorithm that provides a review assignment in the core, thus also establishing the existence of such an assignment, under some conditions. Yet again, the idea of greedy algorithms plays a significant role.

In the last chapter of Part III, we considered a team formation problem, where the goal is to partition a set of $n$ agents into $k$ almost equal-sized teams, when there are friendships among the agents and each agent prefers to be in the same team with as many of her friends as possible. In this work, we defined fairness in two different ways. First, we used the idea of the core to define proportionally fair teams. In particular, as in the previous cases, since there are $n$ agents and the goal is to form $k$ equal-sized teams, each subset of roughly $n/k$ agents is entitled to form its own group. Therefore, a solution is in the core if no subset of agents has an incentive to collectively deviate and form their own group. Moreover, we asked for teams that are envy-free. Here, envy-freeness requires that no agent prefer to exchange teams with another agent, which, in essence, means that no agent prefers to receive the allocation of

another agent, as is commonly seen in social choice theory. For both notions of fairness, we provide approximation guarantees, and in some cases we use greedy algorithms to get these results.

Overall, we see that in this thesis, the idea of the core plays a significant role in many of the chapters. This concept provide essential frameworks for defining and achieving fairness in various scenarios. From clustering to peer review to sortition to team formation, the core serves as a fundamental criterion for ensuring group fairness. As we mentioned in the introduction, the main advantage of the core is that ensures fairness over any possible group and not only to predefined ones. An important aspect that is missed from this thesis though, is that in many cases, the deviation of a group from a given solution does not come without a cost. For example, in the peer-review setting, it is quite non trivial for a community to build its own conference or clear benefits of being part of a larger community should be taken into account before considering isolation. Therefore, it is crucial to incorporate this cost factor into different scenarios. Another pivotal concept in this thesis is the use of greedy algorithms which it is proved to be a practical and effective approach to designing (approximately) fair and efficient algorithms in a wide range of domains.

# Future Vision

In Part I, we focused on optimizing decisions for social good. Throughout this part, social good was measured by the (utilitarian) social welfare, but other approaches can be used in other real-life problems. For example, other social welfare functions or perhaps even notions of ethics (such as fairness or privacy) can serve as yardstick for social good. This can be important in applications involving the allocation of scarce resources, such as blood donation, food donation, or distribution of essential supplies in case of a natural disaster. It is also important to explore the ethical challenges that arise in automated ethical decision-making. For example, how can algorithms for autonomous vehicles use democratic principles to make moral decisions when imposing harm becomes inevitable? I believe that understanding how automated decision-making can be used to maximize social good has tremendous potential for positively transforming the society we live in.

In Part II, we focused on democratic innovations for enhancing civic participation in the current democratic system and unfolded a number of key issues. For example, to ensure that the voice of each citizen is heard, many governmental and non-profit organizations have focused on developing online platforms, through which any citizen can propose their own ideas and other citizens can vote on them. Such a bottom-up approach has many challenges. For example, how can the voting be conducted effectively, given that not every citizen can be asked to vote on every idea? How can all ideas be treated fairly? If citizens have contrasting viewpoints regarding an idea, how can deliberation be used as a mediation tool? It is also important to model and understand deliberation between voters, whether physically or through online platforms, at a mathematical level. We believe that a deliberation process, when designed carefully, can act as the antidote to extreme polarization, which is a serious issue that plagues our society today.

In Part III, we discussed how notions of fairness from economics can be effective in modern decision-making paradigms including machine learning. Moving forward, it is particularly interesting to explore the adaption of other popular fairness notions from the computational social choice literature to machine learning, such as maximin share fairness or group fairness, and exploring other machine learning settings, such as federated learning. It is also interesting to explore how to formally define meaningful notions of

fairness for standalone ML modules (such as a Word2Vec tool), when their future use in an application and its impact on human lives is difficult to predict at the time of the design of the module. Another challenge is for each application and instance to pick the "correct" notion of fairness.As demonstrated in different chapters, different objectives can be incompatible. Therefore, in some cases, it should be taken the decision of which solution is the fairest. We believe thought that there are multiple ethical, legal, philosophical, economic, and many other, considerations that should be taken into account on making this decision. Therefore, researchers with relevant backgrounds can contribute towards this direction and a really interdisciplinary research can give the correct answer. While this exploration contributes towards the ambitious goal of building an overarching theory of algorithmic fairness, this exercise may be futile if the fairness guarantees break down due to strategic manipulations by self-interested users. How do we design systems that are robust to such strategic manipulations? Overall, we believe that ideas of fairness and incentives from economics have the potential to induce the next revolution in computing, and my vision is to help make it a reality.

# Bibliography


H. Abdollahpouri, M. Mansoury, R. Burkeand, and B. Mobasher. The unfairness of popularity bias in recommendation. *arXiv preprint arXiv:1907.13286*, 2019.

A. Abdulkadiroğlu and T. Sönmez. School choice: A mechanism design approach. *American Economic Review*, 93(3):729–747, 2003.

A. Abdulkadiroğlu, P. A. Pathak, A. E. Roth, and T. Sönmez. The boston public school match. *American Economic Review*, 95(2):368–371, 2005.

A. Abera, H. Belay, A. Zewude, B. Gidey, D. Nega, B. Dufera, A. Abebe, T. Endriyas, B. Getachew, H. Birhanu, et al. Establishment of COVID-19 testing laboratory in resource-limited settings: challenges and prospects reported from Ethiopia. *Global Health Action*, 13(1):1841963, 2020.

A. Agarwal, M. Dudík, and Z. S. Wu. Fair regression: Quantitative definitions and reduction-based algorithms. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, volume 97, pages 120–129, 2019.

R. Aharoni, E. C. Milner, and K. Prikry. Unfriendly partitions of a graph. *Journal of Combinatorial Theory, Series B*, 50(1):1–10, 1990.

E. Aigner-Horev and E. Segal-Halevi. Envy-free matchings in bipartite graphs and their applications to fair division. *Information Sciences*, 587:164–187, 2022.

N. Alon, M. Feldman, A. D. Procaccia, and M. Tennenholtz. Strategyproof approximation of the minimax on networks. *Mathematics of Operations Research*, 35(3):513–526, 2010.

G. Amanatidis, G. Birmpas, A. Filos-Ratsikas, A. Hollender, and A. A. Voudouris. Maximum Nash welfare and other stories about EFX. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1758–1765, 2020.

N. Anari, T. Mai, S. O. Gharan, and V. V. Vazirani. Nash social welfare for indivisible items under separable, piecewise-linear concave utilities. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2274–2290, 2018.

J. Angwin, J. Larson, S. Mattu, and L. Kirchner. Machine bias: There's software used across the country to predict future criminals. *And it's biased against blacks. ProPublica*, 23, 2016.

E. M. Arkin, S. W. Bae, A. Efrat, K. Okamoto, J. S. Mitchell, and V. Polishchuk. Geometric stable roommates. *Information Processing Letters*, 109(4):219–224, 2009.

K. Arrow. *Social Choice and Individual Values*. Wiley, 1951.

P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire. The nonstochastic multi-armed bandit problem. *SIAM Journal on Computing*, 32(1):48–77, 2002.

N. Augenblick, J. T. Kolstad, Z. Obermeyer, and A. Wang. Group testing in a pandemic: The role of frequent testing, correlated risk, and machine learning. Working Paper 27457, National Bureau of Economic Research, July 2020. URL `http://www.nber.org/papers/w27457`.

H. Aziz and R. Savani. Hedonic games. In F. Brandt, V. Conitzer, U. Endress, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 15. Cambridge University Press, 2016.

H. Aziz, M. Brill, F. Fischer, P. Harrenstein, J. Lang, and H. G. Seedig. Possible and necessary winners of partial tournaments. *Journal of Artificial Intelligence Research*, 54:493–534, 2015.

H. Aziz, M. Brill, V. Conitzer, E. Elkind, R. Freeman, and T. Walsh. Justified representation in approval-based committee voting. *Social Choice and Welfare*, 48(2):461–485, 2017.

H. Aziz, A. Bogomolnaia, and H. Moulin. Fair mixing: The case of dichotomous preferences. In *Proceedings of the 20th ACM Conference on Economics and Computation (EC)*, pages 753–781, 2019a.

H. Aziz, I. Caragiannis, A. Igarashi, and T. Walsh. Fair allocation of indivisible goods and chores. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 53–59, 2019b.

H. Aziz, R. Freeman, N. Shah, and R. Vaish. Best of both worlds: Ex ante and ex post fairness in resource allocation. *Operations Research*, 2023a.

H. Aziz, B. E. Lee, and S. M. Chu. Proportionally representative clustering. arXiv:2304.13917, 2023b.

A. Backurs, P. Indyk, K. Onak, B. Schieber, A. Vakilian, and T. Wagner. Scalable fair clustering. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, pages 405–413, 2019.

M. L. Balinski and H. P. Young. *Fair Representation: Meeting the Ideal of One Man, One Vote*. Yale University Press, 1982.

A. Ban and N. Linial. Internal partitions of regular graphs. *Journal of Graph Theory*, 83(1):5–18, 2016.

E. Bargiacchi, T. Verstraeten, D. Roijers, A. Nowé, and H. Hasselt. Learning to coordinate with coordination graphs in repeated single-stage multi-agent decision problems. In *Proceedings of the 35thInternational Conference on Machine Learning (ICML)*, pages 482–490, 2018.

N. Barrot and M. Yokoo. Stable and envy-free partitions in hedonic games. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 67–73, 2019.

D. Baumeister, P. Faliszewski, J. Lang, and J. Rothe. Campaigns for lazy voters: truncated ballots. In *Proceedings of the 11th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 577–584, 2012.

R. F. Baumeister, E. Bratslavsky, C. Finkenauer, and K. D. Vohs. Bad is stronger than good. *Review of general psychology*, 5(4):323–370, 2001.

C. Bazgan, Z. Tuza, and D. Vanderpooten. Efficient algorithms for decomposing graphs under degree constraints. *Discrete Applied Mathematics*, 155(8):979–988, 2007.

C. Bazgan, Z. Tuza, and D. Vanderpooten. Satisfactory graph partition, variants, and generalizations. *European Journal of Operational Research*, 206(2):271–280, 2010.

S. Bera, D. Chakrabarty, N. Flores, and M. Negahbani. Fair algorithms for clustering. In *Proceedings of the 32nd Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 4955–4966, 2019.

R. Berk, H. Heidari, S. Jabbari, M. Kearns, and A. Roth. Fairness in criminal justice risk assessments: The state of the art. *Sociological Methods & Research*, 49(3):3–38, 2018.

A. Beynier, Y. Chevaleyre, L. Gourvès, A. Harutyunyan, J. Lesca, N. Maudet, and A. Wilczynski. Local envy-freeness in house allocation problems. *Autonomous Agents and Multi-Agents Systems (AAMAS)*, 33(5):591–627, 2019.

V. Bilò, G. Monaco, and L. Moscardelli. Hedonic games with fixed-size coalitions. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, 2022. 9287–9295.

M. J. Binnicker. Challenges and controversies to testing for COVID-19. *Journal of clinical microbiology*, 58(11):e01695–20, 2020.

A. Biswas and S. Barman. Fair division under cardinality constraints. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 91–97, 2018.

R. A. Bloodgood, J. G. Short, J. M. Jackson, and J. R. Martindale. A change to pass/fail grading in the first two years at one medical school results in improved psychological well-being. *Economics of Education Review*, 84:655–662, 2020.

A. Blumer, A. Ehrenfeucht, D. Haussler, and M. K. Warmuth. Learnability and the vapnik-chervonenkis dimension. *Journal of the Association for Computing Machinery*, 36(4):929–965, 1989.

A. Bogomolnaia and M. Jackson. Stability of hedonic coalition structures. *Games and Economic Behavior*, 38(2):201–230, 2002.

A. Borodin, O. Lev, N. Shah, and T. Strangway. Primarily about primaries. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, pages 1804–1811, 2019.

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.

S. J. Brams, D. M. Kilgour, and M. R. Sanver. A minimax procedure for electing committees. *Public Choice*, 132:401–420, 2007.

F. Brandl, F. Brandt, M. Greger, D. Peters, C. Stricker, and W. Suksompong. Funding public projects: A case for the Nash product rule. *Journal of Mathematical Economics*, 99:102585, 2022.

F. Brandt, V. Conitzer, U. Endress, J. Lang, and A. D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

V. Brault, B. Mallein, and J.-F. Rupprecht. Group testing as a strategy for COVID-19 epidemiological monitoring and community surveillance. *PLoS computational biology*, 17(3):e1008726, 2021.

R. Bredereck, P. Faliszewski, A. Kaczmarczyk, and R. Niedermeier. An experimental view on committees providing justified representation. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 109–115, 2019.

M. Brill, R. Freeman, S. Janson, and M. Lackner. Phragmén's voting methods and justified representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pages 406–413, 2017.

M. Brill, P. Gölz, D. Peters, U. Schmidt-Kraepelin, and K. Wilker. Approval-based apportionment. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, pages 1854–1861, 2020.

A. Brownback. A classroom experiment on effort allocation under relative grading. *Economics of Education Review*, 62:113–128, 2018.

S. Bubeck and R. Eldan. Multi-scale exploration of convex functions and bandit convex optimization. In *Proceedings of the 29thConference on Learning Theory (COLT)*, pages 583–589, 2016.

S. Bubeck, R. Munos, G. Stoltz, and C. Szepesvári. X-armed bandits. *The Journal of Machine Learning Research*, 12:1655—-1695, 2011.

S. Bubeck, Y. T. Lee, and R. Eldan. Kernel-based methods for bandit convex optimization. In *Proceedings of the 49thAnnual Symposium on Theory of Computing (STOC)*, pages 72–85, 2017.

E. Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.

T. Calders, F. Kamiran, and M. Pechenizkiy. Building classifiers with independency constraints. In *Proceedings of the IEEE International Conference on Data Mining Workshops (ICDM)*, pages 13–18, 2009.

J. Cameron and W. D. Pierce. Reinforcement, reward, and intrinsic motivation: A meta-analysis. *Review of Educational research*, 64(3):363–423, 1994.

I. Caragiannis, A. D. Procaccia, and N. Shah. When do noisy votes reveal the truth? *ACM Transactions on Economics and Computation*, 4(3):15, 2016.

I. Caragiannis, X. Chatzigeorgiou, G. A. Krimpas, and A. A. Voudouris. Optimizing positional scoring rules for rank aggregation. *Artificial Intelligence*, 267:58–77, 2019a.

I. Caragiannis, D. Kurokawa, H. Moulin, A. D. Procaccia, N. Shah, and J. Wang. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation*, 7(3): article 12, 2019b.

I. Caragiannis, N. Shah, and A. A. Voudouris. The metric distortion of multiwinner voting. *Artificial Intelligence*, 313:103802, 2022.

M. Chakraborty, K. Y. P. Chua, S. Das, and B. Juba. Coordinated versus decentralized exploration in multi-agent multi-armed bandits. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 164–170, 2017.

L. Charlin and R. Zemel. The toronto paper matching system: an automated paper-reviewer assignment system. In *Proceedings of the ICML Workshop on Peer Reviewing and Publishing Models*, 2013.

J. Chen and S. Roy. Euclidean 3D stable roommates is np-hard. arXiv:2108.03868, 2021.

W. Chen, A. Srivastav, G. Travaglini, et al. *A panorama of discrepancy theory*. Springer, 2014.

X. Chen, B. Fain, C. Lyu, and K. Munagala. Proportionally fair clustering. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, pages 1032–1041, 2019.

Y. Cheng, Z. Jiang, K. Munagala, and K. Wang. Group fairness in committee selection. *ACM Transactions on Economics and Computation (TEAC)*, 8(4):1–18, 2020.

A. Chhabra, K. Masalkovaitė, and P. Mohapatra. An overview of fairness in clustering. *IEEE Access*, 9: 130698–130720, 2021.

F. Chierichetti, R. Kumar, S. Lattanzi, and S. Vassilvitskii. Fair clustering through fairlets. In *Proceedings of the 31st Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 5036–5044, 2017.

L. J. Christensen and K. E. Menzel. The linear relationship between student reports of teacher immediacy behaviors and perceptions of state motivation, and of cognitive, affective, and behavioral learning. *Communication Education*, 47:82–90, 1998.

R. Cole and V. Gkatzelis. Approximating the Nash social welfare with indivisible items. *SIAM Journal on Computing*, 47(3):1211–1236, 2018.

R. Cole, N. Devanur, V. Gkatzelis, K. Jain, T. Mai, V. V. Vazirani, and S. Yazdanbod. Convex program duality, Fisher markets, and Nash social welfare. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pages 459–460, 2017.

L. M. Coleman, L. Jussim, and J. Abraham. Students' r] eactions to teachers' e] valuations: The unique impact of negative feedback 1. *Journal of Applied Social Psychology*, 17(12):1051–1070, 1987.

M. d. 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.

V. Conitzer, R. Freeman, and N. Shah. Fair public decision making. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pages 629–646, 2017.

V. Conitzer, R. Freeman, N. Shah, and J. W. Vaughan. Group fairness for the allocation of indivisible goods. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, pages 1853–1860, 2019.

J. Cox and M. B. Partensky. Spatial localization problem and the circle of Apollonius. arXiv:physics/0701146, 2007.

Á. Cseh, T. Fleiner, and P. Harján. Pareto optimal coalitions of fixed size. *Mechanism and Institution Design*, 4(13):87–108, 2019.

E. Czibor, S. Onderstal, R. Sloof, and M. C. van Praag. Does relative grading help male students? Evidence from a field experiment in the classroom. *Economics of Education Review*, 75:101953, 2020.

M. R. De Georgeo, J. M. De Georgeo, T. M. Egan, K. P. Klee, M. S. Schwemm, H. Bye-Kollbaum, and A. J. Kinser. Containing sars-cov-2 in hospitals facing finite ppe, limited testing, and physical space variability: Navigating resource constrained enhanced traffic control bundling. *Journal of Microbiology, Immunology and Infection*, 54(1):4–11, 2021.

E. de Klerk, M. Laurent, and P. A. Parrilo. A PTAS for the minimization of polynomials of fixed degree over the simplex. *Theoretical Computer Science*, 361(2-3):210–225, 2006.

E. de Klerk, M. Laurent, and Z. Sun. An alternative proof of a PTAS for fixed-degree polynomial optimization over the simplex. *Mathematical Programming*, 151(2):433–457, 2015.

E. L. Deci, R. Koestner, and R. M. Ryan. A meta-analytic review of experiments examining the effects of extrinsic rewards on intrinsic motivation. *Psychological bulletin*, 125(6):627, 1999.

S. Demko and T. P. Hill. Equitable distribution of indivisible objects. *Mathematical Social Sciences*, 16 (2):145–158, 1988.

P. C. Dev. Intrinsic motivation and academic achievement: What does their relationship imply for the classroom teacher? *Remedial and special education*, 18(1):12–19, 1997.

P. Dey and A. Bhattacharyya. Sample complexity for winner prediction in elections. In *Proceedings of the 14th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1421–1430, 2015.

G. N. Dhabaan, W. A. Al-Soneidar, and N. N. Al-Hebshi. Challenges to testing COVID-19 in conflict zones: Yemen as an example. *Journal of global health*, 10(1), 2020.

K. Dhull, S. Jecmen, P. Kothari, and N. B. Shah. Strategyproofing peer assessment via partitioning: The price in terms of evaluators' expertise. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, pages 53–63, 2022.

R. Dorfman. The detection of defective members of large populations. *The Annals of Mathematical Statistics*, 14(4):436–440, 1943.

J. A. Doucette, K. Larson, and R. Cohen. Approximate winner selection in social choice with partial preferences. In *Proceedings of the 12th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2014.

S. Dryden-Peterson, G. E. Velásquez, T. J. Stopka, S. Davey, S. Lockman, and B. O. Ojikutu. Disparities in SARS-CoV-2 testing in massachusetts during the COVID-19 pandemic. *JAMA network open*, 4(2): e2037067–e2037067, 2021.

A. Dubey and A. S. Pentland. Kernel methods for cooperative multi-agent contextual bandits. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, 2020.

P. Dubey and J. Geanakoplos. Grading exams: 100, 99, 98,... or A, B, C? *Games and Economic Behavior*, 69(1):72–94, 2010.

C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. S. Zemel. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference (ITCS)*, pages 214–226, 2012.

F. Echenique and J. Oviedo. A theory of stability in many-to-many matching markets. *Theoretical Economics*, 1(2):233–273, 2006.

E. Eisenberg and D. Gale. Consensus of subjective probabilities: The pari-mutuel method. *The Annals of Mathematical Statistics*, 30(1):165–168, 1959.

E. Elkind and M. Lackner. Structure in dichotomous preferences. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2019–2025, 2015.

E. Elkind, P. Faliszewski, and A. Slinko. Distance rationalization of voting rules. *Social Choice and Welfare*, 45(2):345–377, 2015.

E. Elkind, M. Lackner, and D. Peters. Structured preferences. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 10. 2017.

J. Ely, A. Galeotti, O. Jann, and J. Steiner. Optimal test allocation. *Journal of Economic Theory*, 193: 105236, 2021.

F. Engelstad. The assignment of political office by lot. *Social Science Information*, 28(1):23–50, 1989.

B. Fain, A. Goel, and K. Munagala. The core of the participatory budgeting problem. In *Proceedings of the 12th Conference on Web and Internet Economics (WINE)*, pages 384–399, 2016.

B. Fain, K. Munagala, and N. Shah. Fair allocation of indivisible public goods. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, pages 575–592, 2018.

M. Feldman and Y. Wilf. Strategyproof facility location and the least squares objective. In *Proceedings of the fourteenth ACM conference on Electronic Commerce (EC)*, pages 873–890, 2013.

M. Feldman, S. A. Friedler, J. Moeller, C. Scheidegger, and S. Venkatasubramanian. Certifying and removing disparate impact. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 259–268, 2015.

J. Fishkin. *Democracy When the People Are Thinking: Revitalizing Our Politics Through Public Deliberation*. Oxford University Press, 2018.

G. W. Flake, R. E. Tarjan, and K. Tsioutsiouliklis. Graph clustering and minimum cut trees. *Internet Mathematics*, 1(4):385–408, 2004.

B. Flanigan, P. Gölz, A. Gupta, and A. D. Procaccia. Neutralizing self-selection bias in sampling for sortition. In *Proceedings of the 34th Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 6528–6539, 2020.

B. Flanigan, P. Gölz, A. Gupta, B. Hennig, and A. D. Procaccia. Fair algorithms for selecting citizens' assemblies. *Nature*, 596:548–552, 2021a.

B. Flanigan, G. Kehne, and A. D. Procaccia. Fair sortition made transparent. In *Proceedings of the 35th Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 25720–25731, 2021b.

D. K. Foley. Resource allocation and the public sector. *Yale Economics Essays*, 7:45–98, 1967.

D. Fotakis. Incremental algorithms for facility location and k-median. *Theoretical Computer Science*, 361(2-3):275–313, 2006.

R. Freeman, S. M. Zahedi, and V. Conitzer. Fair and efficient social choice in dynamic settings. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 4580–4587, 2017.

D. Gale and L. S. Shapley. College admissions and the stability of marriage. *Americal Mathematical Monthly*, 69(1):9–15, 1962.

J. Gan, W. Suksompong, and A. A. Voudouris. Envy-freeness in house allocation problems. *Mathematical Social Sciences*, 101:104–106, 2019.

C. Gangl, J. Maly, M. Lackner, and S. Woltran. Aggregating expert opinions in support of medical diagnostic decision-making. In *Proceedings of the 11th International Workshop on Knowledge Representation for Health Care (KR4HC-2019)*, 2019.

M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.

J. Garg and P. McGlaughlin. Improving Nash social welfare approximations. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 294–300, 2019.

J. Garg, M. Hoefer, and K. Mehlhorn. Approximating the Nash social welfare with budget-additive valuations. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2326–2340, 2018.

N. Garg, T. Kavitha, A. Kumar, K. Mehlhorn, and J. . Mestre. Assigning papers to referees. *Algorithmica*, 58(1):119–136, 2010.

J. Gastil and E. O. Wright, editors. *Legislature by Lot: Transformative Designs for Deliberative Governance*. Verso, 2019.

G. George and S. Marvin. *Puzzle-math*. New York, NY, USA: Viking Press, 1958.

M. U. Gerber and D. Kobler. Algorithmic approach to the satisfactory graph partitioning problem. *European Journal of Operational Research*, 125(2):283–291, 2000.

S. Gillen, C. Jung, M. Kearns, and A. Roth. Online learning with an unknown fairness metric. In *Proceedings of the 32nd Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 2600–2609, 2018.

D. B. Gillies. *Some theorems on n-person games*. Princeton University, 1953.

N. Goldberg and G. Rudolf. On the complexity and approximation of the maximum expected value all-or-nothing subset. *Discrete Applied Mathematics*, 283:1–10, 2020.

S. Gollapudi, K. Kollias, and B. Plaut. Almost envy-free repeated matching in two-sided markets. In *Proceedings of the 16th Conference on Web and Internet Economics (WINE)*, pages 3–16, 2020.

C. Gollier and O. Gossner. Group testing against Covid-19. Technical report, EconPol Policy Brief, 2020.

A. Graefe. Accuracy of vote expectation surveys in forecasting elections. *Public Opinion Quarterly*, 78 (S1):204–232, 2014.

T. Gray and J. Bunte. The effect of grades on student performance: Evidence from a quasi-experiment. *College Teaching*, 70(1):15–28, 2022.

A. Gupta, G. Guruganesh, and M. Schmidt. Approximation algorithms for aversion k-clustering via local k-median. In *Proceedings of the 43rd International Colloquium on Automata, Languages and Programming (ICALP)*, pages 66:1–66:13, 2016.

T. R. Guskey. Five obstacles to grading reform. *Educational Leadership*, 69(3):16, 2011.

J. Y. Halpern and D. C. Parkes. Journals for certification, conferences for rapid dissemination. *Communications of the ACM*, 54(8):36–38, 2011.

M. Hardt, E. Price, and N. Srebro. Equality of opportunity in supervised learning. In *Proceedings of the 30th Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 3315–3323, 2016.

T. Hashimoto, M. Srivastava, H. Namkoong, and P. Liang. Fairness without demographics in repeated loss minimization. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, pages 1929–1938, 2018.

J. Hatfield, F. Kojima, Y. Narita, et al. Promoting school competition through school choice: A market design approach. *Stanford: Stanford University Capital and Economic Opportunity Working Group Working Paper*, 18:2011, 2011.

E. Hazan and K. Y. Levy. Bandit convex optimization: Towards tight bounds. In *Proceedings of the 28th Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 784–792, 2014.

N. Hazon, Y. Aumann, S. Kraus, and M. Wooldridge. On the evaluation of election outcomes under uncertainty. *Artificial Intelligence*, 189:1–18, 2012.

S. T. Hedetniemi, D. P. Jacobs, and K. E. Kennedy. A theorem of ore and self-stabilizing algorithms for disjoint minimal dominating sets. *Theoretical Computer Science*, 593:132–138, 2015.

S. Hossain, A. Mladenovic, and N. Shah. Designing fairly fair classifiers via economic fairness notions. In *Proceedings of the 29th International World Wide Web Conference (WWW)*, pages 1559–1569, 2020.

R. W. Irving. An efficient algorithm for the "stable roommates" problem. *Journal of Algorithms*, 6(4): 577–595, 1985.

M. Jones, H. Nguyen, and T. Nguyen. An efficient algorithm for fair multi-agent multi-armed bandit with low regret. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, pages 8159–8167, 2023.

J. Jonnerby, P. Lazos, E. Lock, F. Marmolejo-Cossío, C. B. Ramsey, M. Shukla, and D. Sridhar. Maximising the Benefits of an Acutely Limited Number of COVID-19 Tests. *ArXiv (preprint)*, pages 1–12, 2020a. URL `http://arxiv.org/abs/2004.13650`.

J. Jonnerby, P. Lazos, E. Lock, F. Marmolejo-Cossío, C. B. Ramsey, and D. Sridhar. Test and Contain: A Resource-Optimal Testing Strategy for COVID-19. *AI for Social Good Workshop*, 1, 2020b.

M. Joseph, M. Kearns, J. H. Morgenstern, and A. Roth. Fairness in learning: Classic and contextual bandits. In *Proceedings of the 30th Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 325–333, 2016.

C. Jung, S. Kannan, and N. Lutz. A center in your neighborhood: Fairness in facility location. In *Proceedings of the Symposium on Foundations of Responsible Computing (FORC)*, page 5:1–5:15, 2020.

D. Kahneman and A. Tversky. On the interpretation of intuitive probability: A reply to jonathan cohen. *Cognition*, 7(4):409–411, 1979.

M. Kaneko and K. Nakamura. The nash social welfare function. *Econometrica: Journal of the Econometric Society*, pages 423–435, 1979.

M. M. Kavanagh, N. A. Erondu, O. Tomori, V. J. Dzau, E. A. Okiro, A. Maleche, I. C. Aniebo, U. Rugege, C. B. Holmes, and L. O. Gostin. Access to lifesaving medical resources for african countries: COVID-19 testing and response, ethics, and politics. *The Lancet*, 395(10238):1735–1738, 2020.

M. Kearns, S. Neel, A. Roth, and Z. S. Wu. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, pages 2569–2577, 2018.

B. Klaus and M. Walzl. Stable many-to-many matchings with contracts. *Journal of Mathematical Economics*, 45(7–8):422–434, 2009.

R. Kleinberg, A. Slivkins, and E. Upfal. Multi-armed bandits in metric spaces. In *Proceedings of the 40thAnnual Symposium on Theory of Computing (STOC)*, pages 681–690, 2008.

R. Kleinberg, A. Niculescu-Mizil, and Y. Sharma. Regret bounds for sleeping experts and bandits. *Machine learning*, 80(2-3):245–272, 2010.

R. Kleinberg, A. Slivkins, and E. Upfal. Bandits and experts in metric spaces. *Journal of the ACM (JACM)*, 66(4):1–77, 2019.

A. Kobren, B. Saha, and A. McCallum. Paper matching with local fairness constraints. In *Proceedings of the 25th International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 1247–1257, 2019.

H. Konishi, T. Quint, and J. Wako. On the Shapley-Scarf economy: the case of multiple types of indivisible goods. *Journal of Mathematical Economics*, 35(1):1–15, 2001.

N. Kumar and B. Raichel. Fault tolerant clustering revisited. In *Proceedings of the 25th Canadian Conference on Computational Geometry (CCCG)*, 2013.

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.

M. Lackner and P. Skowron. Approval-based committee voting: Axioms, algorithms, and applications. Technical report, arXiv:2007.01795 [cs.GT], 2021.

M. Lackner and P. Skowron. *Multi-winner voting with approval preferences*. Springer Nature, 2023.

T. L. Lai and H. Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1):4–22, 1985.

H. Landemore. Deliberation, cognitive diversity, and democratic inclusiveness: An epistemic argument for the random selection of representatives. *Synthese*, 190:1209–1231, 2013.

P. Landgren, V. Srivastava, and N. E. Leonard. Distributed cooperative decision-making in multiarmed bandits: Frequentist and bayesian algorithms. *2016 IEEE 55th Conference on Decision and Control (CDC)*, pages 167–172, 2016.

J. Lang, M. S. Pini, F. Rossi, D. Salvagnin, K. B. Venable, and T. Walsh. Winner determination in voting trees with incomplete preferences and weighted votes. *Autonomous Agents and Multi-Agent Systems*, 25(1):130–157, 2012.

D. B. Larremore, B. Wilder, E. Lester, S. Shehata, J. M. Burke, J. A. Hay, M. Tambe, M. J. Mina, and R. Parker. Test sensitivity is secondary to frequency and turnaround time for COVID-19 screening. *Science advances*, 7(1), 2021.

G. P. Latham and E. A. Locke. New developments in and directions for goal-setting research. *European Psychologist*, 12(4):290, 2007.

E. Lee. APX-hardness of maximizing Nash social welfare with indivisible items. *Information Processing Letters*, 122:17 – 20, 2017.

J. Levin and B. Nalebuff. An introduction to vote-counting schemes. *Journal of Economic Perspectives*, 9(1):3–26, 1995.

M. S. Lewis-Beck. Election forecasting: principles and practice. *The British Journal of Politics and International Relations*, 7(2):145–164, 2005.

M. S. Lewis-Beck and C. Tien. Voters as forecasters: a micromodel of election prediction. *International Journal of Forecasting*, 15(2):175–184, 1999.

J. W. Lian, N. Mattei, R. Noble, and T. Walsh. The conference paper assignment problem: Using order weighted averages to assign indivisible goods. In S. A. McIlraith and K. Q. Weinberger, editors, *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18)*, pages 1138–1145, 2018.

K. Liao and D. Guo. A clustering-based approach to the capacitated facility location problem 1. *Transactions in GIS*, 12(3):323–339, 2008.

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.

X. Liu, T. Suel, and N. Memon. A robust model for paper reviewer assignment. In *Proceedings of the 8th ACM Conference on Recommender systems (RecSys)*, pages 25–32, 2014.

Y. Liu, G. Radanovic, C. Dimitrakakis, D. Mandal, and D. C. Parkes. Calibrated fairness in bandits. arXiv preprint arXiv:1707.01875, 2017.

E. Lock, F. J. Marmolejo-Cossío, J. Jonnerby, N. Rajgopal, H. A. Guzmán-Gutiérrez, L. A. Benavides-Vázquez, J. R. Tello-Ayala, and P. Lazos. Optimal testing and containment strategies for universities in Mexico amid COVID-19. In *Proceedings of the 1st Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO)*, pages 1–9. 2021.

E. Lock, F. Marmolejo-Cossío, E. Micha, and A. D. Procaccia. Welfare-maximizing pooled testing. *arXiv preprint arXiv:2206.10660*, 2022.

J. B. Main and B. Ost. The impact of letter grades on student effort, course selection, and major choice: A regression-discontinuity analysis. *The Journal of Economic Education*, 45(1):1–10, 2014.

P. Manurangsi and W. Suksompong. Almost envy-freeness for groups: Improved bounds via discrepancy theory. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 335–341, 2021.

P. J. McEwan, S. Rogers, and A. Weerapana. Grade sensitivity and the economics major at a women's college. In *AEA papers and proceedings*, volume 111, pages 102–06, 2021.

M. McKay and D. Manlove. The three-dimensional stable roommates problem with additively separable preferences. In *Proceedings of the 14th International Symposium on Algorithmic Game Theory (SAGT)*, pages 266—-280, 2021.

N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, and A. Galstyan. A survey on bias and fairness in machine learning. *ACM computing surveys (CSUR)*, 54(6):1–35, 2021.

E. Meskanen and H. Nurmi. Closeness counts in social choice. In *Power, Freedom, and Voting*. Springer-Verlag, 2008.

E. Micha and N. Shah. Proportionally fair clustering revisited. In *47th International Colloquium on Automata, Languages, and Programming (ICALP 2020)*, pages 85:1–85:16, 2020.

D. Mimno and A. McCallum. Expertise modeling for matching papers with reviewers. In *Proceedings of the 13th International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 500–509, 2007.

H. Moulin. *Axioms of Cooperative Decision Making*, volume 15 of *Econometric Society monographs*. Cambridge University Press, 1988.

H. Moulin. *Fair Division and Collective Welfare*. MIT Press, 2003.

A. Narayanan. Translation tutorial: 21 fairness definitions and their politics. In *Proceedings of the 1st Conference on Fairness, Accountability, and Transparency (FAT*)*, 2018.

R. S. Nickerson. What authors want from journal reviewers and editors. *American Psychological*, pages 661—-662, 2005.

J. Nisbet. Adding and averaging grades. *Educational Research*, 17(2):95–100, 1975.

R. O'Dell, M. Wattenhofer, and R. Wattenhofer. The paper assignment problem. *Technical Report/ETH Zurich, Department of Computer Science*, 491, 2005.

M. Olsen, L. Bækgaard, and T. Tambo. On non-trivial nash stable partitions in additive hedonic games with symmetric 0/1-utilities. *Information Processing Letters*, 112(23):903–907, 2012.

C. O'Neil. *Weapons of math destruction: How big data increases inequality and threatens democracy*. Broadway Books, 2016.

O. Ore. Theory of graphs. *American Mathematical Society Colloquium*, 38, 1962.

V. Paredes. Grading system and student effort. *Education Finance and Policy*, 12(1):107–128, 2017.

B. Paria, K. Kandasamy, and B. Póczos. A flexible framework for multi-objective bayesian optimization using random scalarizations. In *Proceedings of the 35th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, 2019.

J. M. Parker. *Randomness and Legitimacy in Selecting Democratic Representatives*. PhD thesis, University of Texas at Austin, 2011.

D. C. Parkes and R. V. Vohra. Algorithmic and economic perspectives on fairness. arXiv:1909.05282, 2019.

V. Patil, G. Ghalme, V. Nair, and Y. Narahari. Achieving fairness in stochastic multi-armed bandits. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5379–5386, 2020.

G. Patro, A. Biswas, N. Ganguly, K. Gummadi, and A. Chakraborty. Fairrec: Two-sided fairness for personalized recommendations in two-sided platforms. In *Proceedings of the 29th International World Wide Web Conference (WWW)*, pages 1194–1204, 2020.

J. Payan and Y. Zick. I will have order! optimizing orders for fair reviewer assignment. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 440–446, 2022.

A. Peng, J. Z. Forde, Y. Shavit, and J. Frankle. Strengthening subcommunities: Towards sustainable growth in ai research. *arXiv preprint arXiv:2204.08377*, 2022.

D. Peters, G. Pierczyński, N. Shah, and P. Skowron. Market-based explanations of collective decisions. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5656–5663, 2021a.

D. Peters, G. Pierczyński, and P. Skowron. Proportional participatory budgeting with additive utilities. In *Proceedings of the 35th Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 12726–12737, 2021b.

G. Pierczyński and P. Skowron. Core-stable committees under restricted domains. In *Proceedings of the 18th International Conference on Web and Internet Economics (WINE)*, pages 311–329, 2022.

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.

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.

F. Pukelsheim. *Proportional Representation: Apportionment Methods and Their Applications*. Springer, 2014.

J. Rawls. A theory of justice. *Cambridge (Mass.)*, 1971.

M. A. Rodriguez and J. Bollen. An algorithm to determine peer-reviewers. In *Proceedings of the 17th ACM Conference on Information and Knowledge Management (CIKM)*, pages 319–328, 2008.

D. Rohe, P. Barrier, M. Clark, D. Cook, K. Vickers, and P. A. Decker. The benefits of pass-fail grading on stress, mood, and group cohesion in medical students. *Mayo Clinic Proceedings*, 81(11):1443–1448, 2006.

C. Rösner and M. Schmidt. Privacy preserving clustering with constraints. In *Proceedings of the 45th International Colloquium on Automata, Languages and Programming (ICALP)*, pages 96:1—96:14, 2018.

A. Roth. Stability and polarization of interests in job matching. *Econometrica: Journal of the Econometric Society*, pages 47–57, 1984.

A. Roth and M. Sotomayor. Two-sided matching. *Handbook of game theory with economic applications*, 1:485–541, 1992.

D. Rothschild and J. Wolfers. Forecasting elections: Voter intentions versus expectations. *SSRN Electronic Journal*, 2011.

S. L. Saidman, A. E. Roth, T. Sönmez, M. U. Ünver, and F. L. Delmonico. Increasing the opportunity of live kidney donation by matching for two and three way exchanges. *Transplantation*, 81:773–782, 2006.

L. Sánchez-Fernández, E. Elkind, M. Lackner, N. Fernández, J. A. Fisteus, P. Basanta Val, and P. Skowron. Proportional justified representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pages 670–676, 2017.

K. H. Shafique. *Partitioning a graph in alliances and its application to data clustering*. PhD thesis, University of Central Florida, 2004.

N. B. Shah. Principled methods to improve peer review. *Technical Report*, 2019.

N. B. Shah. Challenges, experiments, and computational solutions in peer review. *Communications of the ACM*, 65(6):76–87, 2022.

S. Shahrampour, A. Rakhlin, and A. Jadbabaie. Multi-armed bandits in multi-agent networks. In *Proceedings of the 42nd IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2786–2790, 2017.

S. Shalev-Shwartz and S. Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge University Press, 2014.

L. S. Shapley and H. Scarf. On cores and indivisibility. *Journal of Mathematical Economics*, 1(1):23–37, 1974.

V. Shenmaier. The problem of a minimal ball enclosing k points. *Journal of Applied and Industrial Mathematics*, 7(3):444–448, 2013.

A. S. Sikora. Mathematical theory of student assessment through grading. `http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.714.8666&rep=rep1&type=pdf`, 2015. Accessed: 2022-08-15.

P. Skowron, M. Lackner, M. Brill, D. Peters, and E. Elkind. Proportional rankings. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 409–415, 2017.

L. Sless, N. Hazon, S. Kraus, and M. Wooldridge. Forming k coalitions and facilitating relationships in social networks. *Artificial Intelligence*, 259:217–245, 2018.

A. Slivkins et al. Introduction to multi-armed bandits. *Foundations and Trends® in Machine Learning*, 12(1-2):1–286, 2019.

M. A. O. Sotomayor. Three remarks on the many-to-many stable matching problem. *Mathematical Social Sciences*, 38(1):55–70, 1999.

H. Steinhaus. The problem of fair division. *Econometrica*, 16:101–104, 1948.

I. Stelmakh. Towards fair, equitable, and efficient peer review. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, pages 15736–15737, 2021.

I. Stelmakh, N. B. Shah, and A. Singh. Peerreview4all: Fair and accurate reviewer assignment in peer review. In *Algorithmic Learning Theory*, pages 828–856, 2019.

M. Stiebitz. Decomposing graphs under degree constraints. *Journal of Graph Theory*, 23(3):321–324, 1996.

P. Stone. *The Luck of the Draw: The Role of Lotteries in Decision Making*. Oxford University Press, 2011.

T. Sühr, A. Biega, M. Zehlike, K. Gummadi, and A. Chakraborty. Two-sided fairness for repeated matchings in two-sided markets: A case study of a ride-hailing platform. In *Proceedings of the 25th International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 3082–3092, 2019.

K. Tadenuma. Partnership, solidarity, and minimal envy in matching problems. In *Social Ethics and Normative Economics*, pages 155–167. 2011.

C. J. Taylor. On the optimal assignment of conference papers to reviewers. *Technical Report*, 2008.

Z. Terzopoulou, A. Karpov, and S. Obraztsova. Restricted domains of dichotomous preferences with possibly incomplete information. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5726–5733, 2021.

W. R. Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.

H. D. Tran, G. Cabanac, and G. Hubert. Expert suggestion for conference program committees. In *Proceedings in the 11th International Conference on Research Challenges in Information Science (RCIS)*, pages 221–232, 2017.

G. D. L. Travis and H. M. Collins. New light on old boys: Cognitive and institutional particularism in the peer review system. *Science, Technology, & Human Values*, 16(3):322–341, 1991.

A. Tumasjan, T. O. Sprenger, P. G. Sandner, and I. M. Welpe. Predicting elections with twitter: What 140 characters reveal about political sentiment. In *Proceedings of the 4th AAAI conference on Weblogs and Social Media (ICWSM)*, 2010.

University of Waterloo. University of waterloo graduate studies, grading scheme prior to fall 2001. `https://uwaterloo.ca/graduate-studies-academic-calendar/general-information-and-regulations/grades-and-grading`, 2022. Accessed: 2022-08-15.

University of Western Ontario. University of western ontario grading scheme. `https://registrar.uwo.ca/academics/grades_progression_and_graduation.html#gpa`, 2022. Accessed: 2022-08-15.

B. Ustun, Y. Liu, and D. C. Parkes. Fairness without harm: Decoupled classifiers with preference guarantees. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, pages 6373–6382, 2019.

D. Van Reybrouck. *Against Elections: The Case for Democracy*. Random House, 2016.

V. Vapnik and A. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16(2):264—280, 1971.

H. R. Varian. Equity, envy, and efficiency. *Journal of Economic Theory*, 9:63–91, 1974.

Victoria University of Wellington. Victoria university of wellington grading scheme. `https://www.wgtn.ac.nz/students/study/progress/grades`, 2022. Accessed: 2022-08-15.

G. Vrettas and M. Sanderson. Conferences versus journals in computer science. *Journal of the Association for Information Science and Technology*, 66(12):2674–2684, 2015.

M. J. Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.

D. Walther. Picking the winner(s): Forecasting elections in multiparty systems. *Electoral Studies*, 40: 1–13, 2015.

F. Wang, B. Chen, and Z. Miao. A survey on reviewer assignment problem. In *Proceedings of the 21st International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems (IEA/AIE)*, pages 718–727, 2008.

D. H. Wedell, A. Parducci, and D. Roman. Student perceptions of fair grading: A range-frequency analysis. *The American Journal of Psychology*, pages 233–248, 1989.

K. Wojtas and P. Faliszewski. Possible winners in noisy elections. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI)*, pages 1499–1505, 2012.

M. Woodroofe. A one-armed bandit problem with a concomitant variable. *Journal of the American Statistical Association*, 74(368):799–806, 1979.

Q. Wu and A. Roth. The lattice of envy-free matchings. *Games and Economic Behavior*, 109:201–211, 2018.

Y. Xu, X. Zhao, Han Shi, and N. B. Shah. On strategyproof conference peer review. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 616–622, 2019.

Y. Yang. On the tree representations of dichotomous preferences. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 644–650, 2019.

F. Yates. Systematic sampling. *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 241(834):345–377, 1948.

Y. Yokoi. Envy-free matchings with lower quotas. *Algorithmica*, 82(2):188–211, 2020.

H. P. Young. Condorcet's theory of voting. *The American Political Science Review*, 82(4):1231–1244, 1988.

Y. Yue, J. Broder, R. Kleinberg, and T. Joachims. The K-armed dueling bandits problem. *Journal of Computer and System Sciences*, 78(5):1538–1556, 2012.

M. B. Zafar, I. Valera, M. Rodriguez, and K. P. Gummadi. Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. In *Proceedings of the 26th International Conference on World Wide Web (WWW)*, pages 1171–1180, 2017.

M. B. Zafar, I. Valera, M. Gomez-Rodriguez, and K. P. Gummadi. Fairness constraints: A flexible approach for fair classification. *Journal of Machine Learning Research*, 20(75):1–42, 2019.

R. Zemel, Y. Wu, K. Swersky, T. Pitassi, and C. Dwork. Learning fair representations. In *Proceedings of the 30th International Conference on Machine Learning (ICML)*, pages 325–333, 2013.