Approximations for Economic Efficiency and Fairness

by

Kangning Wang

Department of Computer Science
Duke University

Date: ____________________

Approved:

____
Kamesh Munagala, Supervisor

____
Vincent Conitzer

____
Nikolai Gravin

____
Debmalya Panigrahi

Dissertation submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in the Department of Computer Science
in the Graduate School of Duke University
2022

Abstract

Approximations for Economic Efficiency and Fairness

by

Kangning Wang

Department of Computer Science
Duke University

Date: ______________________
Approved:

____
Kamesh Munagala, Supervisor

____
Vincent Conitzer

____
Nikolai Gravin

____
Debmalya Panigrahi

An abstract of a dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Computer Science in the Graduate School of Duke University
2022

Abstract

Efficiency and fairness are two major objectives in designing economic systems – they provide guidelines on the desirability of economic solutions. However, perfectly efficient and fair utopian economies are often nonexistent or impractical, due to incentives, lack of information, computational hardness, etc. In this dissertation, we address this impossibility with methodology from the study of approximation algorithms. Through various cases in the field of computational economics, we demonstrate how to establish and improve the design of provably guaranteed approximations on an ideal yet impossible level of efficiency and fairness. En route, we resolve several open questions in computational social choice and mechanism design.

Acknowledgements

Most importantly, I want to thank my advisor Kamesh Munagala. I am always amazed by his constant dedication to research and education, and his capability in excelling at them. From his exhibition and advice, I learnt keys in various aspects of an academic career; through his care and support, I was able to develop my abilities as a researcher. I owe what I become to him.

For two summers I interned at Google Research. I am very fortunate to have been hosted by Jieming Mao and Renato Paes Leme in 2020, and by Aranyak Mehta in 2021. I have learnt so much through their exhibition of exceptional research taste and technical ability, and have been helped so much by their caring mentorship. For the past two years, I have also had the pleasure of working with Yuan Deng, Guru Guruganesh, Mohammad Mahdian, Vahab Mirrokni, Balu Sivan, Di Wang, Song Zuo from Google. I deeply thank all of them.

Nick Gravin and Zhihao Gavin Tang hosted me for a summer visit in 2019. Thanks to them I had the enjoyment to be with these brilliant and diligent people.

I would like to thank Vince Conitzer, Nick Gravin, and Debmalya Panigrahi for their service on my dissertation committee. Their helpful comments and insights very positively shaped this dissertation.

Thanks to Sid Banerjee as a long-term collaborator, who often provides great new perspectives in our meetings. Thanks to Reza Alijani and Yu Cheng for their mentoring and guidance to research. They made my research life much smoother when I just started at Duke.

During my Ph.D. study, I had the opportunity to work with (then-)undergraduate students Zhihao Jiang, Charles Lyu, Zeyu Shen, and Zhiyi Wang. Their kindness and passion influenced me to stay in academia.

Thanks to all my past collaborators and all who have offered discussion and comments. Many ideas in my research would not have been possible without them.

I would like to thank the very helpful staff members at Duke Computer Science, especially Marilyn Butler, and the warm community of my fellow graduate students: Thanks to Brandon, Chenwei, Hanrui, Jingrong, Jiyao, Keerti, Kevin, Mo, Shao-Heng, Shihan, Xiang, Xiao, Yiheng, Zhengjie, and many others. Their presence makes my Ph.D. life a much more enjoyable experience.

Thanks to my past teachers and coaches, who greatly influenced and cultivated me to become a researcher. They illustrate clearly that education is way beyond a mere selection process.

I want to thank all my friends, those who I constantly bother for company, and those who show up in my need. And, thanks to my parents and extended family members, whose love and encouragement have supported me all these years. To my friends and family I dedicate this dissertation.

Contents

List of Figures

List of Tables

1
Introduction

1.1 Overview

1.1.1 The Holy Grail: Efficiency and Fairness

A primary purpose of the study of economics is to understand how economic structures affect people, and to use that knowledge to improve the well-being of the society. To this end, economists design systems with good *efficiency* and *fairness*.

There are various interpretations of the terms. A commonly adopted, natural view of efficiency is the utilitarian (Benthamite) one, where an economic system is efficient if it maximizes the sum of utilities of its participants. For instance, if a social planner is allocating several unrelated indivisible items to a group of people, then the efficient solution will be to allocate each item to the person who wants it the most.

A prominent criticism of the utilitarian view is that it ignores the distributional inequality: It treats an additional unit of utility for a wealthy, satisfied person the same as that for a person in great need. The pursuit of fairness suggests that maximizing the sum of utilities should not be the sole criterion on the desirability of economic solutions. At one mathematical extreme under the fairness guideline, the max-min (Rawlsian) fairness aims to maximize the utility of the person with the least utility. There are notions that interpolate between the utilitarian view and the max-min fairness, such as the Nash welfare which maximizes the geometric mean of the utilities. Many other fairness concepts exist, such as equal opportunity, balanced representation, envy-freeness, etc; we do not attempt to list all of them.

Note that the utilitarian view needs to add up utilities of different people, bringing the philosophical question of whether utilities of different people are comparable. What we mean by fairness in this dissertation can be expressed as follows: We do not compare utilities of different (groups of) people, but instead, we require an efficiency guarantee to hold for all (groups of) people. It shall become clear as we discuss the models later.

1.1.2 The Impossibility of a Perfect World

For various reasons, we often cannot reach the state of perfect efficiency or fairness as we ideally desire.

In many situations, the lack of information prevents the ability to design an efficient/fair economy. This includes cases where the designers are in want of information, and cases where the participants are so. Imagine a company is making a hiring decision between two candidates. As it cannot perfectly know its compatibility with the candidates, there might be a chance that it makes the wrong decision – leading to an inefficiency.

Different notions of efficiency and fairness themselves often go against each other. As a simple example, utilitarian efficiency and max-min fairness objectives may be optimized at different solutions. We look at another example of methods of charging income tax. Governments necessarily impose taxes to provide public services, as in "nothing is certain except death and taxes". Many countries adopt the structure of progressive income tax, which means the overall tax rate increases as one's income increases, and consider it the fair way of taxation. However, this (along with other common methods such as the flat-rate tax) mis-aligns the incentives: A rational person might output a lower, sub-optimal level of effort as a result of the distortion of marginal income rate per unit of effort. On the other hand, under benign assumptions, the lump-sum tax, which charges a fixed amount from each person, is the only tax form that has a marginal tax rate of 0 and thus leads to efficiency. It is on the extreme opposite of the fairness idea that the rich should be taxed more since they are the ones who benefit the most from public services.

In lots of scenarios, people who participate in the economic systems are rational and self-interested. They act to maximize their own well-being, often neglecting the societal consequences of their actions. The tragedy of the commons is a standard example, where public resources are overused as a consequence of each agent minding their own business. Another example is the Myerson-Satterthwaite impossibility theorem, stating that bilateral trade generally cannot be efficient between two self-interested agents with private information. A desired trade may not happen due to the agents' tendency to lie to get better prices.

In practice, simplicity is desired or even required in designing economic systems. Indeed, if people cannot understand the system, how can they react properly? As an example, many countries use tax brackets for income taxes instead of having a continuous optimized function to describe the marginal tax rates. Additionally, interpretability is needed in some applications such as credit scoring, which also limits the potency of the systems.

Some economic problems are computationally hard. For example, even for the simplistic setting of dividing indivisible items to two people with the same utility function, optimizing the max-min fairness or the Nash welfare is NP-hard by a reduction from the partition problem, which suggests that some instances of the problem cannot be quickly computed. There are computational hardness results about various notions of market equilibria. According a well-known argument, if algorithms cannot quickly find the equilibrium solutions, how can real people find them?

Finally, for some efficiency or fairness notions, desired solutions simply may not exist. For instance, Arrow's impossibility theorem points out that no voting rule can simultaneously satisfy some basic, intuitive axioms.

1.1.3 The Compromise: Guaranteed Approximations

In light of the omnipresent situations where the ideal solutions are unreachable, we take the compromise of providing guaranteed approximations. We use the standard framework in the field of approximation algorithms: First, we define a way to measure efficiency or fairness in an economic system (say higher is better). Next, we define a benchmark that represents the level of efficiency/fairness in an ideal state. We then provide a solution that can guarantee some explicit, uniform fraction of the benchmark in every situation.

In this dissertation, we explore several economic models, mostly in the areas of social choice and mechanism design. We show how to achieve approximate efficiency and/or fairness for each model. All our presented approximations are *constant-approximations*, meaning that for some absolute constant $c > 0$, our solutions can guarantee at least a c-fraction of the benchmark in terms of predefined efficiency/fairness measurements. In the rest of this introductory chapter, we discuss each model that we are about to dive into.

1.2 Approximate Efficiency in Metric-Ordinal Social Choice

Voting in practice usually does not elicit the cardinal utilities of each voter on each candidate. In a single-winner election, social choice rules where they only have access to the ordinal preferences of the candidates are common and extensively studied. An existing model assumes all voters and candidates locate in an unknown metric space, and the cost of a voter on a candidate is the distance between them. The social cost of a candidate is therefore the sum of its distances to all the voters. The efficiency objective naturally aims to minimize the social cost.

This line of work tries to answer the following question: "What social choice rule with only ordinal information achieves the best efficiency guarantee compared with the optimal rule with full information?" In Chapter 2, we describe a simple rule that broke an approximation barrier, which also remains the best till this dissertation among the class of rules with only knowledge of global comparisons of each pair of candidates. We additionally identify structures of the problem, which later helped in settling the optimal efficiency guarantee ratio.

1.3 Approximate Fairness in Committee Selection – the Core

Multi-winner elections are also commonly seen in practice, where voters cast their votes on candidates, with the goal of electing a committee of a given size.

A closely related problem is *participatory budgeting*, where voters use their votes to decide what public projects to fund, conditional on a given limited budget.

In both democratic processes, fairness is an essential objective. A long-standing idea is that every demographic of voters should feel that they have been fairly represented. We study the classical notion of the *core* in these settings, which states that no group of voters can exit the democratic process and find an alternative solution of size scaled down proportionally with their voting power, so that each voter in the deviating group is now better-off.

In Chapter 3, we demonstrate that though the exact core solutions may not exist even in simple situations, approximations of them always do in quite general settings. The approximations are either in the sense that the voters lose a fraction of their voting power during deviation, or in the sense that the voters will only deviate if this makes them better-off by a given multiplier.

1.4 Approximate Efficiency under Chamberlin-Courant Scoring Rule

We again look at multi-winner elections, this time on finding an efficient solution. A classical method in the literature to aggregate voters' preferences is the *(Borda-)Chamberlin-Courant scoring rule*. However, the rule has a major drawback in that its exact optimization is NP-HARD. In Chapter 4, we show that simple election rules can give almost optimal efficiency guarantees through our novel analysis framework, greatly improving the results of existing work.

1.5 Approximate Efficiency in Bilateral Trade

In a bilateral trade, a seller is trying to sell an item to a buyer. The buyer has some utility for getting the item, and the seller has some utility for not selling the item (either from keeping the item or from saving the production cost of it). The utilities are private information and are unknown to each other. We look at the classical Bayesian setup, where the utilities are drawn from publicly known distributions.

In a perfectly efficient world, the item should go to the person with the higher utility. This will maximize the *gains-from-trade* – the expected net utility gain because of the bilateral trade. The Myerson-Satterthwaite impossibility theorem states the unfortunate: Efficiency is generally impossible in bilateral trade between two rational, self-interested people, regardless of the trading mechanism in use, as long as the mechanism itself is not subsidizing the two people in the trade.

Using our approximation framework, we naturally reach the question: Can we always guarantee a constant fraction of the gains-from-trade, compared with that in the perfectly efficient world? In Chapter 5, we answer this in the positive, resolving a major research question in the field.

1.6 Approximate Efficiency and Fairness in Auctions via Persuasion

In a classical Bayesian auction setting, a seller aims to maximize her revenue by selling one item to a group of buyers. The seller has no value/cost on the item, meaning that her producer surplus is equal to his revenue. The buyers as a group, gain a consumer surplus equal to the value of the winner (if there is one) minus their payment.

In an efficient allocation, the item should always go to the buyer with the highest valuation. This may not happen if we simply allow the seller to design the auction, since her objective is different – to maximize the producer surplus. Instead, we use *Bayesian persuasion* to signal to the seller about the buyers' realized values.

If our goal is merely to maximize efficiency, we can directly tell the seller all the information about the realized values. In this case, the auction is efficient, and the seller harvests all the gains-from-trade, leaving the consumer surplus to be 0.

In Chapter 6, we additionally impose the fairness constraint that the mechanism should result in good consumer surplus as well. Following an existing idea, notice that the total surplus is at most the efficient-case gains-from-trade, and the producer surplus is at least her without-signaling revenue, since she can always choose to ignore the signals. Their difference gives a natural upper bound of consumer surplus, which we take as the benchmark. Under commonly used distributional assumptions, we show that a simple signaling scheme can simultaneously approximate the consumer surplus and the producer surplus well, which we consider to be approximately fair to both groups.

1.7 Bibliographic Notes

Chapter 2 is based on joint work [MW19] with Kamesh Munagala.

Chapter 3 is based on joint work [CJMW20, JMW20, MSWW22] with Yu Cheng, Zhihao Jiang, Kamesh Munagala, Yiheng Shen, and Zhiyi Wang.

Chapter 4 is based on joint work [MSW21] with Kamesh Munagala and Zeyu Shen.

Chapter 5 is based on joint work [DMSW22] with Yuan Deng, Jieming Mao and Balasubramanian Sivan.

Chapter 6 is based on joint work [ABMW22] with Reza Alijani, Siddhartha Banerjee, and Kamesh Munagala.

1.7.1 Excluded Publications and Preprints

During my doctoral study, I have also coauthored on the following publications and preprints excluded from this dissertation.

- The work of [CGMW21] is on designing a simple mechanism for revenue approximation in a multi-item auction for a budgeted buyer.
- The work of [MLW22] is on interactive Bayesian persuasion in 2-player games, especially bilateral trade.
- The work of [GTW21] is on designing approximation algorithms for edge-arrival online stochastic bipartite matching.
- The work of [ABG$^+$20] is on designing approximation algorithms for a prophet inequality model with stochastic horizons.
- The work of [DMSW21] is on characterizing optimal pricing schemes for selling to an impatient buyer.
- The work of [GMWW22] is on formulating and proving revenue guarantees of prior-independent auctions for multiple heterogeneous buyers.
- The work of [MMW22] is on approximate regret minimization on a setting where observations are noisy.

- The work of [LFMW21] is on a novel centrality notion to incorporate diversity in social networks.
- The work of [KJW$^+$21] is on core stability in a machine learning problem of classification.

2

Approximate Efficiency in Metric-Ordinal Social Choice

2.1 Introduction

Collective decision making and social choice are the cornerstones of modern democracy. Mathematically, a social choice rule is a function that maps ordinal preferences of voters for alternatives (or candidates) to a single winning alternative. One way to characterize social choice rules is by adopting the axiomatic approach. Exemplary works include [Gib73] and [Sat75], where several natural desirable properties of social choice rules are proposed, and many of them are shown to be incompatible with each other.

An alternative approach is to adopt the utilitarian view, where agents have latent cardinal preferences underlying the ordinal preferences [PR06, CP11, BCH$^+$15]. The goal is to optimize certain social objectives, the most common being the sum of utilities of the voters. We view it as a measurement of efficiency. The term *distortion* of a social choice rule refers to the worst-case ratio between the objective value that the rule achieves and the optimal one, where the worst case is over all cardinal preferences that are consistent with the given ordinal preferences. However, it has been established that many common rules have unbounded distortion [PR06] and even randomized social choice rules have distortion of at least $\Omega\left(\sqrt{n}\right)$, where n is the number of candidates [BCH+15].

These impossibility results necessitate imposing additional structure on the underlying cardinal preferences. One common approach is the spatial model, where voters and candidates are located in a certain space (*e.g.* the Euclidean space) and the distance between a voter and a candidate is part of the cost function of the voter [DH66, EH84, EH90, MIMG99, Sch07]. For instance, for the social issue of building a facility serving a neighborhood, it is natural to assume the cost of an individual is the distance between the her and the facility. As another example, consider a government in the process of making a public policy. If people are viewed to be on the left-right political spectrum, the cost of an individual can be modeled as the distance between her and the policy on the spectrum. In both cases, the voters and candidates are located in a metric space and the costs are the corresponding distances.

When the cardinal costs are assumed to be distances in an underlying metric space, it is possible to achieve constant distortion. The work of [ABE+18] shows that selecting any candidate in the uncovered set [Mou86] — for instance, the Copeland voting rule — guarantees distortion of 5. (Please refer to Section 2.2.4 for description of several common voting rules.) They also show a lower bound of 3 for the distortion for any (deterministic) social choice rule. The exact value of the optimal distortion was an open question — neither the upper bound of 5 nor the lower bound of 3 had been improved. The work of [SE17] shows that the well known voting schemes of using scoring rules, as well as the popular single transferable vote, have super-constant distortion. The work of [GKM17] shows the ranked pairs rule and the Schulze method, both of which work on the weighted tournament graph over candidates, have distortion of at least 5 contrary to previous belief that they could improve the bound of 5. The weighted tournament graph places a directed edge between two candidates with weight equal to the portion of voters who prefer the first candidate over the second. The work of [GKM17] then conjectured any social choice rule that only looks at the weighted tournament graph over candidates cannot beat 5.

2.1.1 Our Results

Our first result in Section 2.3 disproves the conjecture that weighted tournament rules cannot improve the distortion bound below 5. We improve the distortion upper bound to $2+\sqrt{5} \approx 4.236$ using a weighted tournament rule. We achieve the bound by generalizing the notion of uncovered sets to a class of weighted rules. We show 4.236 is indeed the best distortion achievable by this new class of social choice rules. This distortion ratio remains the best known among the class of weighted tournament rules.

This still leaves a gap between the best lower bound of 3 and the upper bound of 4.236 for deterministic social choice rules. In Section 2.4, we propose another social choice rule based on bipartite matchings that generalizes uncovered sets in a different way. In Theorem 2.4.4, we show that if a candidate that satisfies our criterion exists, then choosing this candidate yields distortion of 3. In Section 2.4.3. we present a combinatorial conjecture that implies such a candidate always exists. This conjecture is finally proved later by the breakthrough work of [GHS20]. A different, simpler proof of our conjecture is given by [KK22].

2.1.2 Related Work

We refer to the readers a comprehensive survey of the topic by [AFSV21]. Below we briefly mention some interesting lines of related research with their central problems still being open.

Distortion of randomized social choice rules In addition to deterministic social choices rules, randomized rules have been studied in the metric distortion setting. It is possible that they can achieve better distortion than deterministic ones: Randomly selecting a voter to be the dictator achieves distortion of 3 [FFG16, AP17]. On the other hand, no randomized rule can get distortion better than 2.112 [CR22]. The remaining gap gives a fascinating research direction. It is known that no truthful randomized rule can beat 3 [FFG16] and no weighted tournament rule can beat 3 [GKM17]. We focus on deterministic rules since one might argue that randomized rules are unsatisfying for practical implementation and interpretability, and much literature on social choice has therefore focused on deterministic voting rules.

Other Graph Problems in the Metric Distortion Setting The metric distortion framework is powerful enough to be applied to settings beyond that considered in this work. For example, low-distortion algorithms for the maximum bipartite matching problem and other graph problems in metric spaces are proposed in [AS16a, AS16b]. However, distortion bounds for minimum-cost versions of graph problems remain elusive. The best known results still have polynomial dependence on n [FFZ14, CFF$^+$16]. It is an open prominent question whether better bounds are possible.

2.2 Preliminaries

2.2.1 Social Choice Rules

Let $\mathcal{C} = \{1, 2, \ldots, n\}$ be a set of n candidates and $\mathcal{V} = \{1', 2', \ldots, m'\}$ be a set of m voters. Let $\mathcal{A} = \mathcal{C} \cup \mathcal{V}$ be the set of entities in the voting system. For each voter $v \in \mathcal{V}$, she has a strict preference ordering σ_v on the candidates in $\mathcal{C}$. We call $\boldsymbol{\sigma} = \{\sigma_v\}_{v\in\mathcal{V}}$ the voting profile. If $v \in \mathcal{V}$ prefers $c_1 \in \mathcal{C}$ to $c_2 \in \mathcal{C}$, we write $c_1 \succ_v c_2$

or $c_2 \prec_v c_1$. Let $\mathcal{O}_k$ be the family of strict orderings on k elements. A *social choice rule* is a function

$$f : \bigcup_{n,m=1}^{\infty} (\mathcal{O}_n)^m \to \mathbb{N}$$

that takes m orderings on n candidates and outputs one of the n candidates. The output of f is the winner of the social choice rule.

Throughout the paper, we reserve upper-case letters A, B, C, D, X, Y, Z to denote candidates. AB is the set of voters that prefer A to B. ABC is the set of voters that both prefer A to B and prefer B to C. For any set S, we use $|S|$ to denote the size of S. *E.g.*, $|AX|$ is the number of voters who prefer A to X.

We will need the following definition later.

Definition 2.2.1. For any voter $v \in \mathcal{V}$ and any candidate $X \in \mathcal{C}$, $P_v(X) := \{X\} \cup \{Y \in \mathcal{C} \mid Y \succ_v X\}$ is the set of candidates that v likes at least as much as X. Similarly, $Q_v(X) := \{X\} \cup \{Y \in \mathcal{C} \mid Y \prec_v X\}$ is the set of candidates that v likes at most as much as X.

2.2.2 Metric Distortion

We assume the candidates and the voters are located in the same metric space $(\mathcal{A}, d)$, where $d : \mathcal{A} \times \mathcal{A} \to \mathbb{R}_{\geq 0}$ is the metric (distance function). As a metric, d satisfies the following conditions:

1. Positive definiteness: $\forall i, j \in \mathcal{A}, d(i, j) \geq 0$ and $d(i, j) = 0 \iff i = j$.
2. Symmetry: $\forall i, j \in \mathcal{A}, d(i, j) = d(j, i)$.
3. Triangle inequality: $\forall i, j, k \in \mathcal{A}, d(i, j) + d(j, k) \geq d(i, k)$.

We say a metric d is *consistent* with a voting profile $\boldsymbol{\sigma}$ if

$$\forall v \in \mathcal{V}, \forall A, B \in \mathcal{C} : A \succ_v B \implies d(A, v) \leq d(B, v).$$

Here we allow a voter to have equal distances to different candidates. This is only to eliminate the need of infinitesimal distances in later analysis. Similar results can be derived if we require $d(A, v)$ and $d(B, v)$ to be different in the above definition.

We define the *social cost* of a candidate X with respective to a metric d as:

$$S(X, d) = \sum_{v \in \mathcal{V}} d(v, X)$$

and the *distortion* of a social choice rule f as:

$$\Delta(f) = \sup_{\boldsymbol{\sigma}} \sup_{d \text{ consistent with } \boldsymbol{\sigma}} \frac{S(f(\boldsymbol{\sigma}), d)}{\min_{A \in \mathcal{C}} S(A, d)}.$$

2.2.3 Weighted Tournament Graphs

We can construct a tournament graph on the candidates by looking at their pairwise relationships: Let $\mathcal{C}$ be the set of vertices. For each pair of $A, B \in \mathcal{C}$ ($A \neq B$), draw an edge from A to B if A pairwise beats B: $|AB| = |\{v \in \mathcal{V} : A \succ_v B\}| \geq \frac{m}{2}$ (break ties arbitrarily). The result is a directed graph where there is exactly one directed edge between any pair of vertices. If a social choice rule makes decision only based on the tournament graph, it is called a *tournament rule*, or a $C1$ rule [Fis77].

One way to generalize the class of tournament rules is to further consider the margin of each winning. For each pair of $A, B \in \mathcal{C}$ ($A \neq B$), draw an edge from A to B with weight $\frac{|AB|}{m}$. In the resulted *weighted tournament graph*, there are two directed edges between any pair of vertices. A social choice rule making decision only based on the weighted tournament graph is called a *weighted tournament rule*, or a $C2$ rule[1] [Fis77].

2.2.4 Examples of Social Choice Rules

Now we present some social choice rules from existing literature, along with their distortion bounds.

[1] To be precise, in [Fis77], a social choice rule is $C2$ only if it is not $C1$.

Uncovered The uncovered set of a tournament graph consists of those candidates $A \in \mathcal{C}$ such that for any other candidate B, either there is an edge from A to B, or there is another candidate C so that there are an edge from A to C and an edge from C to B [Mou86]. The uncovered set is always nonempty for any tournament graph. Uncovered is a $C1$ rule that selects the alphabetically smallest candidate from the uncovered set. It has distortion of 5 [ABE+18].

Copeland Copeland selects the candidate that has the most pairwise wins, *i.e.*, the candidate with the largest out-degree in the tournament graph (break ties alphabetically). Copeland is a $C1$ rule as well, and it can be seen as Uncovered with a specific tie-breaking method. It does not improve the 5-distortion of Uncovered [ABE+18].

RankedPairs RankedPairs is a $C2$ rule proposed in [Tid87]. In RankedPairs, the edges of the weighted tournament graph are sorted to have decreasing weights. We start from an empty graph and add one edge at a time (in the order of decreasing weights) if it does not create a cycle. A directed acyclic graph (DAG) is resulted in the end, and the source of the DAG is selected as the winning candidate. It achieves distortion of 3 when the tournament graph has a circumference (size of the largest cycle) of at most 4 [ABE+18]. However, its distortion is at least 5 in general [GKM17].

Schulze Schulze is a $C2$ rule proposed in [Sch11]. For two candidates A and B in the weighted tournament graph, define $p(A, B)$ to be the maximum p so that there is a path from A to B with each edge having weight of at least p. It can be shown that there is a candidate $X \in \mathcal{C}$ so that $p(X, Y) \geq p(Y, X)$ for any other $Y \in \mathcal{C}$. This X is selected as the winning candidate by Schulze. Its distortion is at least

5 [GKM17].

OptimalLP In [GKM17], the authors show that the optimal distortion can be achieved by OptimalLP involving solving linear programs. The distortion of selecting candidate A comparing to selecting B in an instance is given by the following linear program:

$$\begin{aligned} P(A, B, \boldsymbol{\sigma}) := \quad & \text{maximize} \quad && \sum_{v \in \mathcal{V}} d(A, v) \\ & \text{subject to} \quad && \sum_{v \in \mathcal{V}} d(B, v) = 1, \\ & && d \text{ is consistent with } \boldsymbol{\sigma}. \end{aligned}$$

OptimalLP selects the candidate $\arg\min_{A \in \mathcal{C}} \max_{B \in \mathcal{C}} P(A, B, \boldsymbol{\sigma})$. Its distortion was known to be between 3 and 5 (inclusive). Our work implies OptimalLP has distortion of at most $2 + \sqrt{5}$.

2.3 Weighted-Uncovered-Set Voting Rule

In this section, we present the WeightedUncovered rule that has distortion of $2+\sqrt{5} \approx 4.236$ (Theorem 2.3.4), improving the best existing upper bound of 5 for any deterministic social choice rule. WeightedUncovered is based on the λ-*weighted uncovered set* (Definition 2.3.1), a generalization of the uncovered set [Mou86]. We show our analysis is tight by giving a matching lower bound (Theorem 2.3.10). We also give a lower bound for the fairness ratio (Theorem 2.3.12).

We start by introducing λ-*weighted uncovered sets*.

Definition 2.3.1. Let $\lambda \in [0.5, 1]$ be a constant. The λ-*weighted uncovered set* is the set of candidates $A \in \mathcal{C}$ such that for any candidate $B \neq A$, we have:

- either $|AB| \geq (1 - \lambda)m$,
- or there is another node $C \notin \{A, B\}$ so that $|AC| \geq (1 - \lambda)m$ and $|CB| \geq \lambda m$.

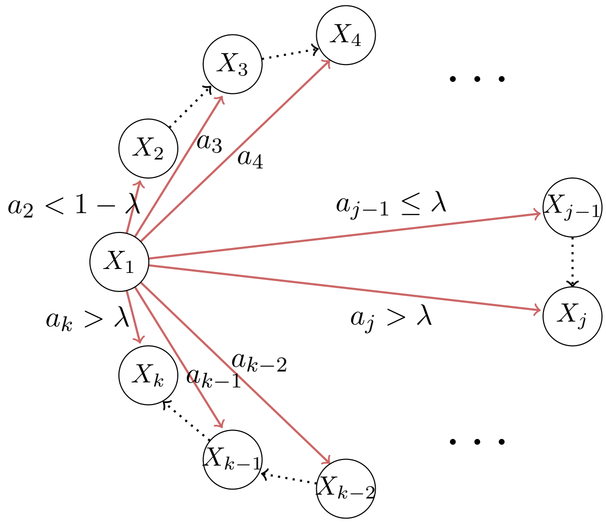

FIGURE 2.1: The Cycle in Lemma 2.3.2

Now we present our key lemma showing the λ-weighted uncovered set is a meaningful concept.

Lemma 2.3.2. *For any $\lambda \in [0.5, 1]$, the λ-weighted uncovered set is nonempty.*

Proof. We prove by contradiction. Suppose there is no such node A. Then there must be nodes $X_1, X_2, \ldots, X_k, X_{k+1} = X_1$ in the graph. If we set $A = X_i$ and $B = X_{i+1}$ ($i = 1, 2, \ldots, k$), then both of the conditions in the lemma are violated. Figure 2.1 shows such a cycle.

Let $a_j = \frac{|X_1 X_j|}{m}$. We have $a_k = \frac{|X_1 X_k|}{m} > \lambda$ and $a_2 = \frac{|X_1 X_2|}{m} < 1 - \lambda$. When j goes from k down to 2, a_j has to drop below λ at some time, so we have $a_j > \lambda$ while $a_{j-1} \leq \lambda$. Thus we have $|X_1 X_j| > \lambda m$ and $|X_{j-1} X_1| \geq (1 - \lambda)m$. This contradicts with the assumption that when $A = X_{j-1}$ and $B = X_j$, there is no such $C = X_1$ satisfying the second condition in the lemma. $\square$

Throughout this section, we use φ to denote the golden ratio $\frac{\sqrt{5}-1}{2} \approx 0.618$.

Definition 2.3.3. WEIGHTEDUNCOVERED is a social choice rule that picks the alphabetically smallest candidate in the φ-weighted uncovered set.

Because of Lemma 2.3.2, WeightedUncovered is well-defined and always picks a candidate in the φ-weighted uncovered set. We state the main property of WeightedUncovered below in Theorem 2.3.4 and prove it later.

Theorem 2.3.4. WeightedUncovered *has distortion of at most* $\left(2+\sqrt{5}\right) \approx 4.236$.

2.3.1 Main Theorem: Improved Distortion Bound

Recall that $S(X,d) = \sum_{v\in\mathcal{V}} d(v,X)$ is the social cost if X is selected and the metric is d. First, if $|AB| \geq (1-\varphi)m$, the social cost gap between choosing A and B can be easily bounded.

Lemma 2.3.5. *If* $|AB| \geq (1-\varphi)m$, *then* $S(A,d) \leq \left(2+\sqrt{5}\right) S(B,d)$ *for any metric* d.

Proof. We have

$$\begin{aligned}
&\left(2+\sqrt{5}\right) S(B,d) - S(A,d)\\
=&\sum_{v\in AB}\left(\left(2+\sqrt{5}\right)d(B,v)-d(A,v)\right)+\sum_{v\in BA}\left(\left(2+\sqrt{5}\right)d(B,v)-d(A,v)\right)\\
\geq&\sum_{v\in AB}\left(1+\sqrt{5}\right)d(B,v)-\sum_{v\in BA}d(A,B)\\
\geq&\sum_{v\in AB}\frac{1+\sqrt{5}}{2}d(A,B)-\sum_{v\in BA}d(A,B)\\
=&d(A,B)\cdot\left(\frac{3+\sqrt{5}}{2}|AB|-m\right)\\
\geq&0.
\end{aligned}$$

□

When showing Copeland has distortion of at most 5, [ABE+18] crucially uses the following lemma:

Lemma 2.3.6 (Theorem 15 in [ABE+18])**.** *If $|AC| \geq \frac{1}{2} \cdot m$ and $|CB| \geq \frac{1}{2} \cdot m$, then $S(A, d) \leq 5S(B, d)$ for any metric d.*

Lemma 2.3.6 itself is tight, but worst-case analysis is deployed when we assume every edge in the unweighted tournament graph has weight of 0.5. Actually, if every edge has weight approaching 0.5, any social choice rule will have distortion approaching 3 by a similar argument as Lemma 2.3.5. Instead of having the stronger condition in the uncovered set that $|AB| \geq 0.5m$, we allow those candidate A's with $|AB| \geq (1-\lambda)m$ to be in the λ-weighted uncovered set. In this way, we can use unbalanced $|AC|$ and $|CB|$ and still guarantee the λ-weighted uncovered set is nonempty (Lemma 2.3.2). It turns out that better distortion can be achieved with carefully chosen parameters.

Lemma 2.3.7. *If $|AC| \geq (1-\varphi)m$ and $|CB| \geq \varphi m$, then $S(A, d) \leq \left(2 + \sqrt{5}\right) S(B, d)$ for any metric d.*

To prove Lemma 2.3.7. We need the following technical lemma from [ABE+18], whose proof is included here for completeness. Recall that $Q_v(A)$ is the set of candidates that v likes at most as much as A.

Lemma 2.3.8 (Lemma 14 in [ABE+18])**.** *If $\sum_{v \in \mathcal{V}} d(v, B) \geq \frac{1}{\gamma} \sum_{v \in BA} \min_{C \in Q_v(A)} d(B, C)$, then $S(A, d) \leq (1+\gamma) S(B, d)$ for any metric d.*

Proof. By definition of $Q_v(A)$ and the triangle inequality,

$$\frac{S(A,d)}{S(B,d)} = \frac{\sum_{v\in\mathcal{V}} d(A,v)}{\sum_{v\in\mathcal{V}} d(B,v)} = \frac{\sum_{v\in AB} d(A,v) + \sum_{v\in BA} d(A,v)}{\sum_{v\in\mathcal{V}} d(B,v)}$$

$$\leq \frac{\sum_{v\in AB} d(B,v) + \sum_{v\in BA}(d(A,v) - d(B,v) + d(B,v))}{\sum_{v\in\mathcal{V}} d(B,v)}$$

$$\leq \frac{\sum_{v\in\mathcal{V}} d(B,v) + \sum_{v\in BA} \min_{C\in Q_v(A)}(d(C,v) - d(B,v))}{\sum_{v\in\mathcal{V}} d(B,v)}$$

$$\leq 1 + \frac{\sum_{v\in BA} \min_{C\in Q_v(A)} d(B,C)}{\sum_{v\in\mathcal{V}} d(B,v)} \leq 1+\gamma. \qquad \square$$

Proof of Lemma 2.3.7. We perform different analysis depending on whether $d(C,B) \geq d(A,B)$ or $d(C,B) < d(A,B)$.

Case (1): When $d(C,B) \geq d(A,B)$:

$$\sum_{v\in\mathcal{V}} d(v,B) \geq \sum_{v\in CB} d(v,B) \geq \frac{1}{2}\sum_{v\in CB} d(C,B) = \frac{1}{2}|CB|\cdot d(C,B).$$

Using the assumption that $d(C,B) \geq d(A,B)$ in this case,

$$\sum_{v\in\mathcal{V}} d(v,B) \geq \frac{\varphi m}{2}\cdot d(A,B) \geq \frac{\varphi}{2}\cdot \sum_{v\in BA} d(A,B).$$

Then we apply Lemma 2.3.8 to get $S(A,d) \leq \left(1+\frac{2}{\varphi}\right) S(B,d) = \left(2+\sqrt{5}\right) S(B,d)$.

Case (2): When $d(C,B) < d(A,B)$:

$$\sum_{v\in\mathcal{V}} d(v,B) \geq \sum_{v\in AB\cup BAC\cup CBA} d(v,B)$$

$$\geq \frac{1}{2}\sum_{v\in AB} d(A,B) + \sum_{v\in BAC} d(v,B) + \frac{1}{2}\sum_{v\in CBA} d(B,C).$$

Notice that for $v\in BAC$, $d(v,B)+d(B,C) \geq d(v,C) \geq d(v,A) \geq d(A,B) - d(v,B)$.

Thus,

$$\sum_{v\in\mathcal{V}} d(v,B) \geq \frac{1}{2}\sum_{v\in AB} d(A,B) + \frac{1}{2}\sum_{v\in BAC}(d(A,B)-d(B,C)) + \frac{1}{2}\sum_{v\in CBA} d(B,C)$$

$$= \frac{1}{2}\left((|AB|+|BAC|)\cdot d(A,B) + (|CBA|-|BAC|)\cdot d(B,C)\right)$$

$$= \frac{1}{2}\left((|AB|+|BAC|)\cdot (d(A,B)-d(B,C)) + (|CBA|+|AB|)\cdot d(B,C)\right).$$

Because $AC \subseteq BAC \cup AB$ and $CB \subseteq CBA \cup AB$,

$$\sum_{v\in\mathcal{V}} d(v,B) \geq \frac{1}{2}\left(|AC|\cdot(d(A,B)-d(B,C)) + |CB|\cdot d(B,C)\right)$$

$$\geq \frac{1}{2}\left((1-\varphi)m\cdot(d(A,B)-d(B,C)) + \varphi m\cdot d(B,C)\right).$$

The second inequality above comes from the conditions that $|AC| \geq (1-\varphi)m$, $|CB| \geq \varphi m$ and the assumption that $d(A,B) > d(B,C)$ in this case. Therefore,

$$\sum_{v\in\mathcal{V}} d(v,B)$$

$$\geq \frac{1}{2}\left(\frac{1-\varphi}{\varphi}\cdot|CA|\cdot(d(A,B)-d(B,C)) + \varphi\cdot|BA|\cdot d(B,C)\right)$$

$$\geq \frac{\varphi}{2}\left((|CBA|+|BCA|)\cdot(d(A,B)-d(B,C)) + (|CBA|+|BCA|+|BAC|)\cdot d(B,C)\right)$$

$$= \frac{\varphi}{2}\left((|CBA|+|BCA|)\cdot d(A,B) + |BAC|\cdot d(B,C)\right)$$

$$= \frac{\varphi}{2}\left(\sum_{v\in CBA\cup BCA} d(A,B) + \sum_{v\in BAC} d(B,C)\right),$$

where the second inequality is from $CBA \cup BCA \subseteq CA$ and $CBA \cup BCA \cup BAC = BA$. Again by $BA = CBA \cup BCA \cup BAC$, we apply Lemma 2.3.8 to get $S(A,d) \leq \left(1+\frac{2}{\varphi}\right) S(B,d) = \left(2+\sqrt{5}\right) S(B,d)$. □

Proof of Theorem 2.3.4. By Lemma 2.3.2, WeightedUncovered always selects a candidate in the uncovered set. By Lemma 2.3.5, Lemma 2.3.7 and the definition of WeightedUncovered, WeightedUncovered has distortion of at most $2 + \sqrt{5}$. $\square$

2.3.2 Lower Bounds

An immediate question is whether our upper bound is tight. Also, is it possible to choose a $\lambda \neq \varphi$ in the λ-weighted uncovered set to improve the distortion? In this section, we show that our choice of $\lambda = \varphi$ and our analysis are indeed tight. Before showing the lower bounds, we complete the definition of λ-weighted uncovered sets for $\lambda < 0.5$.

Definition 2.3.9. Let $\lambda \in [0, 0.5)$ be a constant. The λ-*weighted uncovered set* is the set of candidates $A \in \mathcal{C}$ such that for any candidate $B \neq A$, we have:

- either $|AB| \geq \lambda m$ (*different from the case where* $\lambda \geq 0.5$),
- or there is another node $C \notin \{A, B\}$ so that $|AC| \geq (1-\lambda)m$ and $|CB| \geq \lambda m$.

We did not define λ-weighted uncovered sets for $\lambda < 0.5$ earlier because for any such λ, selecting any candidate in the λ-weighted uncovered set can have distortion worse than 5. Now we show lower bounds for all $\lambda \in [0, 1]$.

Theorem 2.3.10. *For any $\lambda \in [0, 1]$, choosing the alphabetically smallest candidate in the λ-weighted uncovered set has distortion of at least $2 + \sqrt{5}$.*

Proof. The tight examples are illustrated in Figure 2.2. In both cases, all points lie on a straight line. In the left case, let p portion of voters be of type v_1 and $1 - p$ portion of voters be of type v_2.[2] When $p \leq 1 - \varphi$, the distortion of choosing A is at least $\frac{(1-\varphi)\cdot 1+\varphi\cdot 2}{(1-\varphi)\cdot 1+\varphi\cdot 0} = 2 + \sqrt{5}$. Since A is in the p-weighted uncovered set and $(1-p)$-weighted uncovered set, we rule out $\lambda \leq 1 - \varphi$ and $\lambda \geq \varphi$.

[2] Technically this is only doable for $\lambda \in \mathbb{Q}$, but we can add an arbitrarily small portion of other voters to make A in the λ-weighted uncovered set.

When $\lambda \in (1-\varphi, \varphi)$, consider the right case in Figure 2.2. Let $1-\lambda$ portion of voters be of type v_1 and λ portion of voters be of type v_2. A is in the λ-weighted uncovered set. However, the distortion for choosing A comparing to B is at least $\frac{(1-\lambda)\cdot 2+\lambda\cdot 3}{(1-\lambda)\cdot 0+\lambda\cdot 1} = 3 + 2\cdot\frac{1-\lambda}{\lambda}$. When $\lambda < \varphi$, the distortion is greater than $2+\sqrt{5}$. $\square$

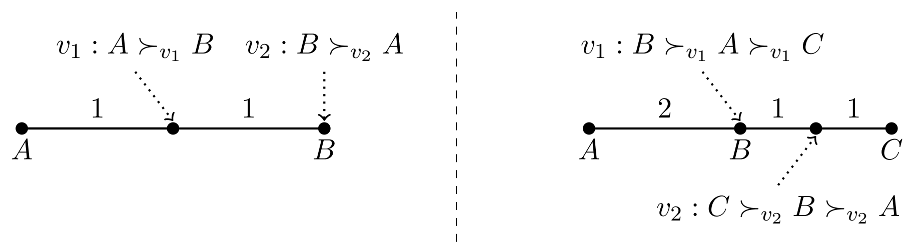

FIGURE 2.2: Tight Examples for WEIGHTEDUNCOVERED with Other Parameters

The *fairness ratio* of social choice rules was proposed in [GKM17]:

Definition 2.3.11. The *fairness ratio* $\Phi(f)$ of a social choice rule f is:

$$\Phi(f) = \sup_{\boldsymbol{\sigma}} \sup_{d \text{ consistent with } \boldsymbol{\sigma}} \max_{1\le k\le m} \frac{\max_{R\subseteq\mathcal{V}:|R|=k} d(v, f(\boldsymbol{\sigma}))}{\min_{A\in\mathcal{C}} \max_{R\subseteq\mathcal{V}:|R|=k} d(v, A)}.$$

The fairness ratio $\Phi(f)$ is at least the distortion $\Delta(f)$ because it approximates the sum of k largest costs for any k including m. The distortion of COPELAND is 5 [ABE+18] and its fairness ratio is 5 [GKM17] as well. Since WEIGHTEDUNCOVERED improves the distortion, one may ask whether it improves the fairness ratio as well. We answer this in the negative:

Theorem 2.3.12. *For any $\lambda \in [0,1]$, choosing the alphabetically smallest candidate in the λ-weighted uncovered set has a fairness ratio of at least 5.*

Proof. When $\lambda = 0$ or $\lambda = 1$, the uncovered set contains every candidate so the fairness ratio is $+\infty$. Otherwise, consider a setting where $\mathcal{C} = \{A, B, C\}$ and $\mathcal{V}$ contains multiple copies of v_1 and v_2. The distances between pairs of points are given in Table 2.1. One can verify they satisfy the triangle inequality.

The preference of v_1 is $C \succ B \succ A$ and that of v_2 is $C \succ A \succ B$. Let λm voters be of v_1 type and $(1-\lambda)m$ voters be of v_2 type.[3] A is in the λ-weighted uncovered set, but the fairness ratio for choosing A is at least 5 comparing to B, given by the worst-off 1 voter (v_1). $\square$

Table 2.1: Adjacency Matrix of the Tight Example for Fairness

Points	A	B	C	v_1	v_2
A	0	4	4	**5**	**3**
B	4	0	2	**1**	**1**
C	4	2	0	1	3
v_1	5	1	1	0	2
v_2	3	1	3	2	0

2.4 Matching-Uncovered-Set Voting Rule

In this section, we present the MatchingUncovered rule, which is based on an alternative generalization of uncovered sets. We introduce *matching uncovered sets* and the MatchingUncovered rule in Section 2.4.1, after which we analyze its distortion (Theorem 2.4.4) in Section 2.4.2. In Section 2.4.3, we propose a conjecture (Conjecture 2.4.8) which implies the distortion of MatchingUncovered is 3. We use computer-assisted search to verify it holds when at most 14 entities are involved. This conjecture was finally proved true by [GHS20, KK22].

2.4.1 Description

Recall that as in Definition 2.2.1, $P_v(A)$ is the set of candidates that v likes as least as much as A and $Q_v(A)$ is the set of candidates that v likes as most as much as A.

[3] Again, this is only doable for $\lambda \in \mathbb{Q}$, but we can add an arbitrarily small portion of other voters to make the λ-weighted uncovered set include A. Instead of looking at the worst-off 1 voter in the end, we look at the worst-off ε portion of voters for a small ε.

Definition 2.4.1. The *matching uncovered set* is the set of candidates $A \in \mathcal{C}$ such that: For any other candidate $B \neq A$, there is a perfect matching in the bipartite graph $G(A, B)$ constructed in the following way:

- The vertices on the left and vertices on the right are both $\mathcal{V}$.
- For a vertex v on the left and a vertex v' on the right, we draw an edge between them if $P_v(B) \cap Q_{v'}(A) \neq \varnothing$.

Example 2.4.2. Let $\mathcal{C} = \{A, B, C\}$ and $\mathcal{V} = \{1, 2, 3\}$. The preference profile is:

$$A \succ_1 B \succ_1 C$$
$$B \succ_2 C \succ_2 A$$
$$C \succ_3 A \succ_3 B$$

A is in the matching uncovered set because $G(A, B)$, $G(A, C)$ both have perfect matchings, as shown in Figure 2.3 and Figure 2.4. By cyclic symmetry, B and C are in the matching uncovered set too.

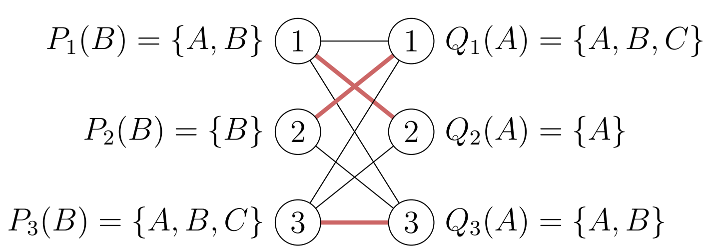

FIGURE 2.3: The Bipartite Graph $G(A, B)$

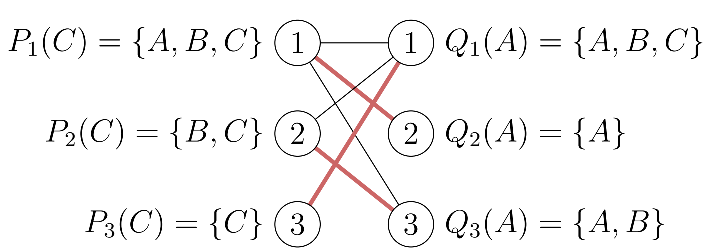

FIGURE 2.4: The Bipartite Graph $G(A, C)$

Definition 2.4.3. MatchingUncovered is a social choice rule that picks the alphabetically smallest candidate in the matching uncovered set. It picks the alphabetically smallest candidate if the matching uncovered set is empty.

2.4.2 Distortion Analysis

The next theorem is the main reason why we introduce matching uncovered sets.

Theorem 2.4.4. *Selecting any candidate from the matching uncovered set guarantees distortion of at most* 3*.*

Corollary 2.4.5. *If the matching uncovered set is nonempty,* MatchingUncovered *has distortion of at most* 3*.*

Before proving Theorem 2.4.4, we introduce a new notation specifically for proving distortion of 3. Fixing the preference orderings $\{\sigma_k\}_{k=1}^m$, define

$$U(v, A, B, d) := d(A, v) - 3d(B, v).$$

Lemma 2.4.6. *A social choice rule f has distortion of at most 3 if and only if for any voting profile $\boldsymbol{\sigma} = \{\sigma_k\}_{k=1}^m$ and for any other candidate $B \neq f(\boldsymbol{\sigma})$,*

$$\sum_{v \in \mathcal{V}} U(v, f(\boldsymbol{\sigma}), B, d) \leq 0$$

for any metric d that is consistent with $\boldsymbol{\sigma}$.

Proof. Let $A = f(\boldsymbol{\sigma})$. A social choice rule f has distortion of at most 3 if and only if for any profile $\boldsymbol{\sigma}$ and any metric d consistent with $\boldsymbol{\sigma}$, $\sum_{v \in \mathcal{V}} d(A, v) \leq 3 \sum_{v \in \mathcal{V}} d(B, v)$ for any other $B \in \mathcal{C}$, *i.e.*, $\sum_{v \in \mathcal{V}} U(v, A, B, d) \leq 0$. □

Lemma 2.4.7. *For a voter v and any consistent metric d:*

1. *If v has preference $Y \succ_v B \succ_v A \succ_v X$, then $U(v, A, B, d) \leq d(B, X) - d(B, Y)$.*

2. *If v has preference $Y \succ_v B \succ_v A$, then $U(v, A, B, d) \leq d(A, B) - d(B, Y)$.*

3. *If v has preference $B \succ_v A \succ_v X$, then $U(v, A, B, d) \leq d(B, X)$.*

4. *If v has preference $A \succ_v B$, then $U(v, A, B, d) \leq -d(A, B)$.*

5. *For any v, $U(v, A, B, d) \leq d(A, B)$.*

Proof. By the preference of v and triangle inequality:

1. If v has preference $Y \succ_v B \succ_v A \succ_v X$, then
$$\begin{aligned} U(v, A, B, d) &= (d(v, A) - d(v, B)) - (d(v, B) + d(v, B)) \\ &\leq (d(v, X) - d(v, B)) - (d(v, B) + d(v, Y)) \\ &\leq d(B, X) - d(B, Y). \end{aligned}$$

2. If v has preference $Y \succ_v B \succ_v A$, then
$$\begin{aligned} U(v, A, B, d) &= (d(v, A) - d(v, B)) - (d(v, B) + d(v, B)) \\ &\leq (d(v, A) - d(v, B)) - (d(v, B) + d(v, Y)) \\ &\leq d(A, B) - d(B, Y). \end{aligned}$$

3. If v has preference $B \succ_v A \succ_v X$, then
$$\begin{aligned} U(v, A, B, d) &= (d(v, A) - d(v, B)) - 2d(v, B) \\ &\leq d(v, X) - d(v, B) \\ &\leq d(B, X). \end{aligned}$$

4. If v has preference $A \succ_v B$, then
$$\begin{aligned} U(v, A, B, d) &= 2(d(v, A) - d(v, B)) - (d(v, A) + d(v, B)) \\ &\leq -(d(v, A) + d(v, B)) \\ &\leq -d(A, B). \end{aligned}$$

5. For any v,

$$\begin{aligned} U(v, A, B, d) &= d(v, A) - d(v, B) - 2d(v, B) \\ &\leq d(v, A) - d(v, B) \\ &\leq d(A, B). \end{aligned}$$ □

Proof of Theorem 2.4.4. Using Lemma 2.4.6, we only need to show $\sum_{v \in \mathcal{V}} U(v, A, B, d) \leq 0$ if there is a perfect matching in $G(A, B)$.

We divide the set of voters into two disjoint subsets: AB (those who prefer A to B) and BA. By (4) in Lemma 2.4.7, we have

$$\begin{aligned} \sum_{v \in \mathcal{V}} U(v, A, B, d) &= \sum_{v \in AB} U(v, A, B, d) + \sum_{v \in BA} U(v, A, B, d) \\ &\leq -|AB| \cdot d(A, B) + \sum_{v \in BA} U(v, A, B, d). \end{aligned}$$

For a voter v, we use L_v to denote the vertex on the left representing $P_v(B)$ and R_v the vertex on the right representing $Q_v(A)$. Because there is a perfect matching in the graph $G(A, B)$, there has to be a matching of size at least $m - 2|AB|$ within the vertices representing BA: $\{L_v \mid v \in BA\} \cup \{R_v \mid v \in BA\}$. We write $M(u)$ to denote the boolean variable whether the vertex u is matched in this smaller matching. If $M(u)$ is true, we write $K(u)$ to denote the common candidate that u uses in the matching. Note that the common candidate cannot be A or B between vertices representing BA. We have

$$
\begin{aligned}
&\sum_{v\in BA} U(v,A,B,d)\\
&= \sum_{v\in BA:M(L_v)\wedge M(R_v)} U(v,A,B,d) + \sum_{v\in BA:M(L_v)\wedge\neg M(R_v)} U(v,A,B,d)\\
&\quad + \sum_{v\in BA:\neg M(L_v)\wedge M(R_v)} U(v,A,B,d) + \sum_{v\in BA:\neg M(L_v)\wedge\neg M(R_v)} U(v,A,B,d)\\
&\le \sum_{v\in BA:M(L_v)\wedge M(R_v)} (d(B,K(R_v)) - d(B,K(L_v))) + \sum_{v\in BA:M(L_v)\wedge\neg M(R_v)} (d(A,B) - d(B,K(L_v)))\\
&\quad + \sum_{v\in BA:\neg M(L_v)\wedge M(R_v)} d(B,K(R_v)) + \sum_{v\in BA:\neg M(L_v)\wedge\neg M(R_v)} d(A,B)\\
&= \sum_{v\in BA:M(L_v)\wedge\neg M(R_v)} d(A,B) + \sum_{v\in BA:\neg M(L_v)\wedge\neg M(R_v)} d(A,B)\\
&= \sum_{v\in BA:\neg M(R_v)} d(A,B).
\end{aligned}
$$

The inequality above is by Lemma 2.4.7. The second last step is because the terms of $d(B,K(R_v))$ and $-d(B,K(L_v))$ cancel out in a matching. Therefore,

$$
\begin{aligned}
\sum_{v\in\mathcal{V}} U(v,A,B,d) &\le -|AB|\cdot d(A,B) + \sum_{v\in BA} U(v,A,B,d)\\
&\le -|AB|\cdot d(A,B) + |\{v\in BA \mid \neg M(R_v)\}|\cdot d(A,B)\\
&\le 0.
\end{aligned}
$$

The last step is because $|\{v\in BA \mid \neg M(R_v)\}|$ is the number of unmatched vertices on the right in the smaller matching within BA, which is at most $|BA|-(m-2|AB|) = |AB|$. □

2.4.3 Existence Conjecture

We present a combinatorial conjecture which, if true, will guarantee MatchingUncovered has distortion of 3.

Conjecture 2.4.8. *For any $\mathcal{C} = \{X_1, X_2, \ldots, X_n\}$ and $\mathcal{V}$, for any voting profile, at least one of the bipartite graphs $G(X_1, X_2), G(X_2, X_3), G(X_3, X_4), \ldots, G(X_{n-1}, X_n)$ and $G(X_n, X_1)$ has a perfect matching.*

Conjecture 2.4.8 immediately implies that MatchingUncovered has distortion of 3: If the matching uncovered set is empty, we can start at any vertex A, find the graph $G(A, B)$ without a perfect matching, then go to B and repeat until we find a cycle $X_1, X_2, \ldots, X_t$. We delete any candidate not on the cycle from the voting profile and maintain relative orders for other candidates. Since $G(X_k, X_{k+1})$ does not have a perfect matching before the deletions, $G(X_k, X_{k+1})$ constructed after the deletions does not have a perfect matching either. However, Conjecture 2.4.8 denies the existence of such a cycle. ($\mathcal{C}$ in Conjecture 2.4.8 is the set of candidates on the cycle.)

Our conjecture was later proved by [GHS20, KK22], showing that MatchingUncovered indeed achieves a distortion of 3.

Theorem 2.4.9 ([GHS20, KK22])**.** *Conjecture 2.4.8 holds, and thus the optimal distortion of 3 can be achieved by a deterministic social choice rule.*

3

Approximate Fairness in Committee Selection – the Core

3.1 Introduction

Consider the following scenario: A city is deciding what public projects to fund using its limited budget b. There is a list of m candidate projects (forming a set C), where each $j \in C$ is associated with a cost s_j. The city needs to select a subset O of these projects whose total cost is at most the budget, that is, $\sum_{j \in O} s_j \leq b$. There are n residents or *voters* (forming a set V) in the city, and each of them has preferences on how the city should spend its budget. These preferences may not be perfectly compatible with each other: For instance, families with children may prefer a public school, while others may prefer a park; people living in the east may prefer projects there, and so for those living in the west. It is desirable to have a fair process to decide on the projects to pay for. In participatory budgeting, voters express their preferences through their votes and influence the decision process. The paradigm has been implemented in numerous cities across the world [Cab04, AS21, pbs, GKSA19].

A similar problem is seen in multiwinner elections [ABES19, End17, BCE$^+$16, CC83, Thi95], where voters select a committee of size b from m candidates. Each voter holds her own opinions on the committees, and a fair method to incorporate the preferences of all voters is called for. Mathematically, it is a special case of the participatory budgeting problem, where each candidate has the same "cost".

In these settings, simple methods to aggregate the preferences may have drawbacks. In majority voting, the utilities of a coherent minority group might be entirely ignored in favor of the majority. Utilitarian ways (maximizing the sum of utilities) may overly focus on a certain group, disregarding the welfare of the vast majority. Indeed, many practically implemented voting rules for choosing parliaments and civic body members, such as the well-known Single Transferable Voting (STV) [TR00], attempt to address precisely this issue. We naturally ask:

> What is a fair solution for participatory budgeting and multiwinner elections, and how do we reach such a solution?

3.1.1 The Core and Its Multiplicative Approximations

Recall that there are n voters forming a set V, and m candidates forming a set C, where candidate j has size (or cost) s_j. We need to choose a subset O of candidates with total size at most b (that is, $\sum_{j \in O} s_j \leq b$). Following social choice terminology, this subset of candidates is called a *committee*, and the problem of choosing a budget-constrained committee is called the *committee selection* problem. Denote the utility of voter i for a committee T by a *utility function* $u_i(T)$. We naturally assume this function is non-negative and monotone, with $u_i(\varnothing) = 0$.

Although there are copious notions of fairness for committee selection, the *core* is a classic and influential one among them. This idea has existed for more than a century [Dro81, Thi95, Lin58], and serves as a strong notion of proportional representation. Towards defining this concept, imagine we split the size b among the voters, so that each voter has an *endowment* of $\frac{b}{n}$ that they can use to "buy" candidates.

A candidate of size s_j requires an endowment of s_j to "buy". A committee $O \subseteq C$ with total size at most b is said to be in the core, if no subset S of voters can deviate and purchase another committee $T \subseteq C$ by pooling their endowments, so that each voter in S prefers the new committee T to the original one O. Note that the total endowment of S is $|S|\frac{b}{n}$, so that this set of voters can buy a committee T of size at most $|S|\frac{b}{n}$. Formally,

Definition 3.1.1 (Core)**.** A committee O is in the *core* if there is no $S \subseteq V$ and committee $T \subseteq C$ with $\sum_{j\in T} s_j \leq \frac{|S|}{n} \cdot b$, such that $u_i(T) > u_i(O)$ for every $i \in S$.

The core has a "fair taxation" interpretation [Lin58, Fol70]. The quantity $\frac{b}{n}$ can be thought of as the tax contribution of a voter, and a committee in the core has the property that no sub-group of voters could have spent their share of tax money in a way that *all* of them were better off. As such it subsumes notions of fairness such as Pareto-optimality, proportionality, and various forms of justified representation [ABC$^+$17, FEL$^+$17, AEH$^+$18] that have been extensively studied in multi-winner election and fairness literature.

Despite the satisfying properties of the core, its strength is also its limitation: Even in the simple setting of unit sizes, integer budget, and additive utilities (the so-called *approval-set* setting with general utilities), the core can be empty. (See for example, [FMS18].)

A natural approach to circumvent this problem is to show the existence of a committee that *multiplicatively* approximates the core. There are two natural directions of introducing a multiplier. First, stemming from the idea where deviation can be costly, we say a deviating voter can only retain a constant fraction of its original endowment.

Definition 3.1.2 (α-Endowment Core)**.** An outcome O is in the α*-endowment core* if there is no $S \subseteq V$ and $T \subseteq C$ with $\sum_{j\in T} s_j \leq \frac{1}{\alpha} \cdot \frac{|S|}{n} \cdot b$, such that $u_i(T) > u_i(O)$

for every $i \in S$.

Another generalization is to say any voter must obtain a constant times more utility in order to be willing to deviate. We define the α-utility core as follows.

Definition 3.1.3 (α-Utility Core)**.** A committee O is in the α*-utility core* if there is no $S \subseteq V$ and $T \subseteq C$ with $\sum_{j \in T} s_j \leq \frac{|S|}{n} \cdot b$, such that $u_i(T) > \alpha \cdot u_i(O \cup \{q\})$ for every $i \in S$ and $q \in C$. We call the candidate q an *additament*.

Note that we introduce the additament in the definition, since no multiplicative approximation is possible without it even in the setting with unit candidate sizes and additive utilities. This follows from examples in previous work [FMS18].

The idea of a bicriteria (multiplicative and additive) approximation to utilities in this context was first introduced by [FMS18]. [PPS21] present an almost identical definition as Definition 3.1.3, except that the additament q must come from the set T. This makes their definition more restrictive and the α-core smaller. They show that when utilities are additive, an $O\left(\log \frac{u_{\max}}{u_{\min}}\right)$-utility core solution not only exists, but can also be computed in polynomial time, where $u_{\max}$ and $u_{\min}$ are the largest and smallest non-zero utilities any voter has for any feasible committee.

3.1.2 Our Results

Our main result is to show the existence of these approximate cores. In Section 3.3, we show the existence of an $O(1)$-endowment core solution as long as the utility functions are monotone. In Section 3.4, we show the existence of an $O(1)$-utility core solution when the utility functions of the voters are monotone and submodular, which can be computed in polynomial time. Submodular utilities capture the notion of "diminishing returns", and are therefore well-motivated in participatory budgeting and multiwinner election settings where each additional project or candidate provides diminishing marginal utility.

Theorem 3.1.4 (Main Theorem for Endowment Approximation)**.** *For general monotone utilities, a* 32*-endowment core is always non-empty.*

Theorem 3.1.5 (Main Theorem for Utility Approximation)**.** *For monotone submodular utilities, a* 67.37*-utility core is always non-empty. One such solution can be computed in polynomial time.*

In Section A.2, we improve the approximation ratio for the well-studied special case of additive utilities for $O(1)$-utility cores:

Theorem 3.1.6 (Additive Utilities)**.** *For additive utilities, a* 9.27*-core is always non-empty.*

Lower Bounds. In Section 3.5, we present our lower bounds. We briefly discuss them here. For the α-endowment core, a natural question is whether a multiplicative approximation factor is necessary. We present the following lower bound to answer it.

Theorem 3.1.7 (Endowment Core Lower Bound)**.** *In the unweighted* Ranking *or additive-utility setting,* $(2-\varepsilon)$*-endowment core of integral size limit* b *may not exist for any constant* $\varepsilon > 0$*.*

The next natural question is whether our results can be extended to arbitrary monotone utilities. We show that this is not possible. We present an example in Section 3.5 to show that an $O(1)$-utility core (or even any $f(n,m)$-utility core where n is the number of voters and m is the number of candidates) may not exist for general monotone utilities.

Theorem 3.1.8 (General Utility Lower Bound)**.** *For general monotone utilities, for any function* $\varphi : \mathbb{Z}^+ \times \mathbb{Z}^+ \to \mathbb{R}^+$*, a* $\varphi(n,m)$*-utility core can be empty.*

We also show that the case of submodular utilities does need a multiplicative approximation to the core, and the 1-utility core can be empty. This justifies the form of Theorem 3.1.5 that involves both a multiplicative approximation and an additament.

Theorem 3.1.9 (Submodular Lower Bound)**.** *For monotone submodular utilities, a* 1.015*-utility core can be empty.*

3.1.3 Technical Overview

Our main technical idea is to consider *randomized versions* or *fractional relaxations* of the problem. For the fractional relaxations, for submodular utilities, we extend the utility to fractional values via the so-called multilinear extension [Von08, CCPV11]; for additive utilities, the extension is straightforward. We then construct an exact or approximate core outcome for the randomized or fractional version. We finally round this solution to find an integer deterministic approximate core solution. The randomized or fractional solution that we use are entirely different for Theorem 3.1.4, Theorem 3.1.5, and Theorem 3.1.6; and we use both dependent rounding [GKPS06, BPR$^+$17] and independent rounding based on the situations. The details are deferred to each section.

During the rounding procedures, we lose the guarantee that any subset of voters will not deviate to a new affordable solution. Instead, we can only guarantee it for any subset of the class of "satisfied" voters. The remaining unsatisfied set of voters can still deviate to a different coalition. To resolve this issue, we propose an iterative rounding framework: We recurse on the unsatisfied voters to find a randomized/fractional solution with a smaller budget on total size. The key idea is that every voter is satisfied at some level of the recursion. We now use the approximation property of the randomized/fractional solution at each level of recursion to argue that the number of voters deviating to any fixed committee is bounded. Summing this argument over all levels completes the proof.

3.1.4 Related Work

Proportionality and the Core. One classic objective in committee selection is achieving fairness via proportionality, where different demographic slices of voters feel they have been represented fairly. This general idea dates back more than a century [Dro81], and has recently received significant attention [CC83, Mon95, BKS07, ABC$^+$17, FEL$^+$17, AEH$^+$18]. In fact, there are several elections, both at a group level and a national level, that attempt to find committees (or parliaments) that provide approximately proportional representation. For instance, the popular Single Transferable Vote (STV) rule is used in parliamentary elections in Ireland and Australia, and in several municipal elections in the USA. This rule attempts to find a proportional solution.

A long line of recent literature has studied the complexity and axiomatization of voting rules that achieve proportionality; see [ABES19, End17, BCE$^+$16] for recent surveys. Proportionality in committee selection arises in many other applications outside of social choice as well. For example, consider a shared cache for data items in a multi-tenant cloud system, where each data item is used by several long-running applications [KFMB17, FGPS19]. Each data item can be treated as a candidate, and each application as a voter whose utility for an item corresponds to the speedup obtained by caching that item. In this context, a desirable caching policy provides proportional speedup to all applications.

The core represents the ultimate form of proportionality: Every demographic of voters feel that they have been fairly represented and do not have the incentive to deviate and choose their own committee of proportionally smaller size which gives all of them higher utility. In the typical setting where these demographic slices are not known upfront, the notion of core attempts to be fair to all subsets of voters.

Fisher Markets. Our fractional solutions are superficially related to the Fisher market equilibrium [BS05, Nas50, AD54] when divisible items need to be *allocated* to agents, and agents' utilities are additive. For the Fisher market, the optimum Nash Welfare solution finds market clearing prices. However, in a Fisher market, the prices are common to the agents while the allocations are different, while in a Lindahl equilibrium, the prices are per-voter while the allocation (or committee) is common and provides shared utility to all the voters. This is a key difference – the Fisher market has a polynomial time algorithm via convex programming [EG59], while no polynomial time algorithm is known for the Lindahl equilibrium even when candidates have unit sizes and voters' utilities are additive (or linear till the maximum size of the candidate). Similarly, though the Nash Welfare solution finds market clearing prices for the Fisher market via strong duality, in the case of public goods, there is no obvious way to interpret the dual of the Nash Welfare solution as market clearing prices. Moreover, for submodular utilities and multilinear extensions, the Nash Welfare objective is no longer a convex program, so that strong duality does not apply.

Approximate Core. In this chapter, we partially focus on approximating the utility voters obtain on deviating (see Definition 3.1.3). As mentioned before, this notion first appeared in [FMS18], and the notion of a single additament in approximation is due to [PPS21]. The latter work present a logarithmic approximation for the special case of additive utilities. [PS20] show that the well-known Proportional Approval Voting method [Thi95] achieves a 2-utility core for the special case where the utilities are additive and candidates are unit size, with each voter having utility either zero or one for each candidate. This algorithm can be viewed as maximizing a discrete version of Nash Welfare, and in essence, Theorem 3.4.4 extends this result to the case of submodular utilities and general costs, showing that it yields a 2-utility core for the *fractional* case of multilinear extension via a polynomial time local search algorithm. The work of [CFLM19] presents a constant approximation for the K-clustering problem, where the committee is a set of K centers in a metric space, and the cost of a voter is the distance to the closest center. However, these ideas do not extend to the committee selection problem we consider in this chapter.

Rounding Techniques. The notion of multilinear extension and correlation gap has been widely used in stochastic optimization [ADSY10], mechanism design [Yan11], and rounding [Von08, CCPV11, CVZ14]. Typically, it has been used to develop computationally efficient approaches; on the other hand, we demonstrate an application to showing a purely existential result. Similarly, rounding of market clearing solutions has been used to show approximately fair allocations of indivisible goods among agents [BKV18, CG15]. The structure of these problems (common prices but different allocations) is very different from ours (common allocations and different prices), and we need different techniques. Again, in contrast with the resource allocation literature, we need the rounding just to show an existence result as opposed to a computational one.

3.2 Model and Preliminaries

Recall that C is a set of m candidates and V is a set of n voters. Each candidate $j \in C$ is associated with a size $s_j > 0$. For each $i \in V$ and each $T \subseteq C$, we denote by $u_i(T)$ the utility of voter i on committee T. We always assume u_i is a monotone set function with $u_i(\varnothing) = 0$. For some of our results, we assume u_i is either submodular or additive. Below are their formal definitions.

- Monotonicity: $u_i(T) \leq u_i(T') \;\; \forall T \subseteq T' \subseteq C$.
- Submodularity: $u_i(T \cup \{t\}) - u_i(T) \geq u_i(T' \cup \{t\}) - u_i(T') \;\; \forall T \subseteq T' \subseteq C, t \notin T'$.

- Additivity: $u_i(T) = \sum_{t \in T} u_i(\{t\}) \ \ \forall T \subseteq C$.

We have a constraint of b on the total size of the committee. We call a committee $O \subseteq C$ feasible (with respect to b) if: $\sum_{j \in O} s_j \le b$.

Our goal is to show the existence of the α-endowment or α-utility core for $\alpha = O(1)$. Recall their definitions:

Definition 3.1.2 (α-Endowment Core)**.** An outcome O is in the *α-endowment core* if there is no $S \subseteq V$ and $T \subseteq C$ with $\sum_{j \in T} s_j \le \frac{1}{\alpha} \cdot \frac{|S|}{n} \cdot b$, such that $u_i(T) > u_i(O)$ for every $i \in S$.

Definition 3.1.3 (α-Utility Core)**.** A committee O is in the *α-utility core* if there is no $S \subseteq V$ and $T \subseteq C$ with $\sum_{j \in T} s_j \le \frac{|S|}{n} \cdot b$, such that $u_i(T) > \alpha \cdot u_i(O \cup \{q\})$ for every $i \in S$ and $q \in C$. We call the candidate q an *additament*.

Randomized Allocations. We consider randomized committees, which we also call *lotteries*, as an intermediate step in our proof later. Let Δ denote a distribution of outcomes, where each $O \in \text{supp}(\Delta)$ is a feasible outcome (i.e. $\text{Cost}(O) \le b$).

Definition 3.2.1 (α-Endowment Stable Lottery)**.** A randomized outcome Δ is an *α-endowment stable lottery* if for any $T \subseteq C$ of size b', we have $\mathbf{E}_{O \sim \Delta}[V(T, O)] < \alpha \cdot \frac{b'}{b} \cdot n$, where $V(T, O)$ denotes the number of voter i's with $u_i(T) > u_i(O)$.

When Δ is deterministic, the definition is equivalent to Definition 3.1.2.

Fractional Allocations. In the sequel, we will consider continuous extensions of the committee and the utility function. We define a *fractional committee* as a m-dimensional vector $\mathbf{x} \ge 0$. The quantity x_j denotes the fraction to which candidate j is allocated.

We denote by $\text{Cost}(\mathbf{x}) := \sum_{j \in C} x_j \cdot s_j$ as the cost of the allocation $\mathbf{x}$. Without loss of generality, we assume that $\sum_{j \in C} s_j > b$.

Similarly, we consider continuous utilities $U_i(\mathbf{x})$ for the voters, whose construction will be specified later. Given this notation, an α-utility fractional core, for $\alpha \geq 1$ is defined as follows:

Definition 3.2.2 (α-Utility Fractional Core)**.** A fractional committee $\mathbf{x} \geq 0$ with $\text{Cost}(\mathbf{x}) \leq b$ lies in the *α-utility fractional core* if there is no $S \subseteq V$ and allocation $\mathbf{z} \geq 0$ with $\text{Cost}(\mathbf{z}) \leq \frac{|S|}{|V|} \cdot b$ such that $U_i(\mathbf{z}) > \alpha U_i(\mathbf{x})$ for all $i \in S$.

When $\alpha = 1$, we call the fractional committee $\mathbf{x}$ as simply lying in the *fractional core*.

3.3 Existence of a 32-Endowment Core Solution

3.3.1 Existence of 2-Endowment Stable Lotteries

In this section, we consider choosing a stable lottery over committees of weight at most b. Recall the definition in this setting.

Definition 3.2.1 (α-Endowment Stable Lottery)**.** A randomized outcome Δ is an *α-endowment stable lottery* if for any $T \subseteq C$ of size b', we have $\mathbf{E}_{O\sim\Delta}[V(T,O)] < \alpha \cdot \frac{b'}{b} \cdot n$, where $V(T,O)$ denotes the number of voter i's with $u_i(T) > u_i(O)$.

We consider the dual formulation of selecting stable lotteries. The existence of a c-approximately stable lottery is equivalent to deciding:

$$\min_{\Delta} \max_{S'} \mathop{\mathbf{E}}_{S\sim\Delta} \left[V(S,S') - c \cdot \frac{w(S')}{b} \cdot n \right] < 0, \tag{3.1}$$

where Δ is a distribution (lottery) over committees of weight at most b. Viewing Δ as a mixed strategy over the "defending" committees and S' as the "attacking" strategy, we treat Eq. (3.1) as a zero-sum game. Duality (or the *min-max principle*) now allows us to swap the order of actions by allowing the attacker to use a mixed strategy. Eq. (3.1) is thus equivalent to

$$\max_{\Delta_a} \min_{S_d : w(S_d) \leq b} \mathop{\mathbf{E}}_{S_a \sim \Delta_a} \left[V(S_d, S_a) - c \cdot \frac{w(S_a)}{b} \cdot n \right] < 0, \tag{3.2}$$

where Δ_a is a lottery over committees of weight at most b chosen by the attacker. This dual view provides a convenient tool for showing the existence of approximately stable lotteries. The rest of the section is devoted to proving Theorem 3.3.1, that we restate here.

Theorem 3.3.1. *For any value b and* all *monotone preferences, a 2-approximately stable lottery over committees of weight at most b always exists.*

Per-Voter Guarantee

Assume we are given a lottery Δ_a. If there is a committee S_a in Δ_a with $w(S_a) > \frac{b}{2}$, then $2 \cdot \frac{w(S_a)}{b} \cdot n > n \geq V(S_d, S_a)$ for any S_d. This implies $V(S_d, S_a) - 2 \cdot \frac{w(S_a)}{b} \cdot n < 0$. Therefore, the attacker can remove these strategies from its lottery, and we can assume Δ_a only has committees with weight at most $\frac{b}{2}$.

Given any distribution Δ_a over committees of weight at most $\frac{b}{2}$, we need to show there is a defending committee S_d with weight at most b, such that

$$\mathop{\mathbf{E}}_{S_a \sim \Delta_a} \left[V(S_d, S_a) - 2 \cdot \frac{w(S_a)}{b} \cdot n \right] < 0.$$

Suppose that the strategy Δ_a chooses S_1 with probability α_1, committee S_2 with probability α_2, ..., S_t with α_t, where $t = 2^{|\mathcal{C}|}$. Let

$$\beta = \mathop{\mathbf{E}}_{S_a \sim \Delta_a} \left[\frac{w(S_a)}{b} \right] = \frac{\sum_{i=1}^{t} \alpha_i \cdot w(S_i)}{b}$$

be the ratio between the expected total weight of the attacking strategy and b, the allowable weight for the defending strategy. We need to find an S_d of weight at most

b so that:

$$\mathop{\mathbf{E}}_{S_a \sim \Delta_a} [V(S_d, S_a)] < 2\beta n. \tag{3.3}$$

We will construct a distribution Δ_d over committees S_d that satisfies a stronger property:

$$\Pr_{S_d \sim \Delta_d, S_a \sim \Delta_a} [S_a \succ_v S_d] < 2\beta \qquad \forall v \in [n]. \tag{3.4}$$

Summing over all voters v implies the existence of a lottery Δ_d satisfying Eq. (3.3), and hence a deterministic committee S_d satisfying the same. This will imply the theorem statement.

Dependent Rounding

Let $p_i = \min\left(1, \frac{\alpha_i}{2\beta}\right)$ for $i \in [t]$. We have:

- $p_i \in [0, 1]$ for all $i \in [t]$; and
- $\sum_{i \in [t]} p_i w(S_i) \leq \frac{1}{2\beta} \sum_{i \in [t]} \alpha_i w(S_i) = \frac{\beta b}{2\beta} = \frac{b}{2}$.

We will construct the defending committee by including each attacker committee $S_i, i \in [t]$ with probability p_i; the details of which are below. We use the random variable X_i to denote whether we include S_i in our defending committee S_d, so that $S_d = \bigcup_{i \in [t]: X_i > 0} S_i$. We therefore obtain a distribution Δ_d over committees S_d.

We use the dependent rounding procedure in [BSS18] to construct $\{X_i\}$ given the $\{p_i\}$. This has the following properties.

- (Almost-Integrality) For any realization of X_i's, all but at most one of them takes value in $\{0, 1\}$ (the remaining one takes value in $[0, 1]$).
- (Correct Marginals) $\mathbf{E}[X_i] = p_i$ for all $i \in [t]$.
- (Preserved Weight) $\Pr\left[\sum_{i \in [t]} w(S_i) \cdot X_i \leq \frac{b}{2}\right] = 1$.

- (Negative Correlation) $\forall T \subseteq [t]$, $\mathbf{E}\left[\prod_{i\in T}(1-X_i)\right] \le \prod_{i\in T}\mathbf{E}[1-X_i] = \prod_{i\in T}(1-p_i)$.

To have full integrality instead of the almost-integrality, in any realization, we include S_i in our S_d as long as $X_i > 0$ (instead of only fully including it when $X_i = 1$). Since we assumed $w(S_i) \le \frac{b}{2}$ for all $i \in [t]$, using the almost-integrality and preserved-weight conditions, for any realization of $\{X_i\}$, the weight of the resulting S_d satisfies

$$w(S_d) \le \sum_{i\in[t]:X_i=1} w(S_i) + \sum_{i\in[t]:X_i\in(0,1)} w(S_i) \le \frac{b}{2} + \frac{b}{2} = b.$$

Analysis

Fix a voter v. W.l.o.g. assume her preference over the sets in Δ_a are $S_1 \succeq_v S_2 \succeq_v \cdots \succeq_v S_t$.

$$\begin{aligned}
\Pr_{S_d\sim\Delta_d, S_a\sim\Delta_a}[S_a \succ_v S_d] &\le \sum_{i\in[t]} \Pr_{S_a\sim\Delta_a}[S_a = S_i \text{ and } X_1 = X_2 = \cdots = X_i = 0] \\
&= \sum_{i\in[t]} \alpha_i \cdot \Pr_{S_a\sim\Delta_a}[X_1 = X_2 = \cdots = X_i = 0] \\
&\le \sum_{i\in[t]} \alpha_i \cdot \prod_{j=1}^{i}(1-p_j) \\
&= \sum_{i\in[t]} 2\beta \cdot p_i \cdot \prod_{j=1}^{i}(1-p_j) < 2\beta.
\end{aligned}$$

Here, the first step follows because when the adversary chooses set S_i, it only beats S_d if S_d included none of $S_1, S_2, \ldots, S_i$. Here, we are using the monotonicity of the preference structure: Since $S_j \succeq_v S_i$ for $j \le i$, this implies $S_d \succeq_v S_i$ when $S_j \subseteq S_d$. The second step follows since the realization of the adversary's lottery is independent of that of the defender, and since $\alpha_i = \Pr[S_a = S_i]$. The third step follows by the negative correlation property of $\{X_j\}$. To see the fourth step, note that if $\alpha_i \geq 2\beta$, then $p_i = 1$, so that $\alpha_i \cdot \prod_{j=1}^{i}(1-p_j) = 0 = 2\beta \cdot p_i \cdot \prod_{j=1}^{i}(1-p_j)$. Otherwise, $\alpha_i = 2\beta p_i$.

To see the final inequality, note that $\sum_{i\in[t]} p_i \prod_{j<i}(1-p_j)$ is the probability of the following stopping process picking some set: Pick S_1 with probability p_1; if not, pick S_2 with probability p_2, and so on. Therefore, $\sum_{i\in[t]} p_i \prod_{j<i}(1-p_j) \leq 1$. Thus $\sum_{i\in[t]} p_i \prod_{j\leq i}(1-p_j) < 1$. This proves Eq. (3.4), and hence Theorem 3.3.1.

3.3.2 Existence of a Deterministic Solution

In this section, we show that a $O(1)$-endowment core solution always exists. We show this by iteratively rounding the lottery constructed in Section 3.3.1. We first restate Theorem 3.1.4.

Definition 3.1.2 (α-Endowment Core)**.** An outcome O is in the *α-endowment core* if there is no $S \subseteq V$ and $T \subseteq C$ with $\sum_{j\in T} s_j \leq \frac{1}{\alpha} \cdot \frac{|S|}{n} \cdot b$, such that $u_i(T) > u_i(O)$ for every $i \in S$.

For the proof, fix a deviating committee S_a of weight $w(S_a)$. Suppose our final committee is T of weight at most b. Our goal is to show that: $V(T, S_a) < 32 \cdot \frac{w(S_a)}{b} \cdot n$.

Good and Bad Committees

Throughout the proof, we fix two constants $0 < \beta \leq \alpha < 1$, whose choice will be determined at the very end. To begin, we define a subroutine that returns a 2-endowment stable lottery (via Theorem 3.3.1) for any subset of voters, and any committee size.

Definition 3.3.2. Given candidate set $[m]$, voter set $\mathcal{V}'$, and committee size b', let $\textsc{Lottery}(\mathcal{V}', b')$ return a lottery Δ over committees of weight at most b' that is 2-endowment stable for the set of voters $\mathcal{V}'$. Similarly, let $V_{\mathcal{V}'}(S, S_a) = |\{v \in \mathcal{V}' \mid S_a \succ_v S\}|$.

Let x_i be the probability that Δ includes committee S_i.

Definition 3.3.3. Given a voter v, we define the set of *good* and *bad* committees relative to Δ, $\mathcal{G}_v(\Delta)$ and $\mathcal{B}_v(\Delta)$ respectively, as follows:

$$\mathcal{G}_v(\Delta) = \left\{ S \subseteq \mathcal{C} \;\middle|\; \sum_{S_i \succeq_v S} x_i \leq 1 - \beta \right\} \qquad \text{and} \qquad \mathcal{B}_v(\Delta) = \left\{ S \subseteq \mathcal{C} \;\middle|\; \sum_{S_i \preceq_v S} x_i \leq \beta \right\}.$$

The idea is that the good committees appear sufficiently high up in v's ranking, while the bad committees are lower down in the ranking. The notion of high and low is relative to the probability mass Δ. The following lemma is immediate.

Lemma 3.3.4. *If $S \notin \mathcal{B}_v(\Delta)$, then $S_a \succ_v S$ only if $S_a \in \mathcal{G}_v(\Delta)$.*

Proof. If $S \notin \mathcal{B}_v(\Delta)$, then $\sum_{S_i \preceq_v S} x_i > \beta$. This implies $\sum_{S_i \succ_v S} x_i < 1 - \beta$. Since $S_a \succ_v S$, we have $\sum_{S_i \succeq_v S_a} x_i < 1 - \beta$, so that $S_a \in \mathcal{G}_v(\Delta)$. $\square$

The next lemma implies that (a) Any committee S_a cannot lie in too many good sets relative to its weight; and (b) There is some committee (with non-zero support in Δ) that does not lie in more than a constant fraction of the bad sets. The previous lemma rules out the possibility where too many voters prefer S_a to such a committee, relative to the weight of S_a, which will be crucial for the algorithm we subsequently design.

Lemma 3.3.5. *Given $\Delta = \textsc{Lottery}(\mathcal{V}', b')$, we have the following upper and lower bounds:*

1. *For all committees S_a, we have*

$$|\{v \in \mathcal{V}' \mid S_a \in \mathcal{G}_v(\Delta)\}| < \frac{2}{\beta} \cdot \frac{w(S_a)}{b'} \cdot |\mathcal{V}'|.$$

2. *There exists S with non-zero support in Δ such that*

$$|\{v \in \mathcal{V}' \mid S \notin \mathcal{B}_v(\Delta)\}| \geq (1 - \beta) \cdot |\mathcal{V}'|.$$

Proof. To see the first part, for any committee S_a, we have: $1 - \sum_{S_i \succeq_v S_a} x_i = \Pr_{S_i \sim \Delta}[S_i \prec_v S_a]$. Summing over $v \in \mathcal{V}'$,

$$\sum_{v \in \mathcal{V}'} \left(1 - \sum_{S_i \succeq_v S_a} x_i\right) = \mathop{\mathbf{E}}_{S_i \sim \Delta} \left[V_{\mathcal{V}'}(S_i, S_a)\right] < \frac{2w(S_a)}{b'} \cdot |\mathcal{V}'|,$$

where the inequality comes from the fact that Δ is 2-endowment stable. Thus there are fewer than $\frac{2}{\beta} \cdot \frac{w(S_a)}{b'} \cdot |\mathcal{V}'|$ voters $v \in \mathcal{V}'$ with $\sum_{S_i \succeq_v S_a} x_i \leq 1 - \beta$, which is necessary for $S_a \in \mathcal{G}_v(\Delta)$.

To see the second part, suppose $S \sim \Delta$. Then for each $v \in \mathcal{V}'$, since $\sum_{S_i \notin \mathcal{B}_v(\Delta)} x_i \geq 1 - \beta$, we have: $\Pr\left[S \notin \mathcal{B}_v(\Delta)\right] \geq 1 - \beta$. Therefore, the expected number (over the choice $S \sim \Delta$) of v such that $S \notin \mathcal{B}_v(\Delta)$ is at least $(1 - \beta) \cdot |\mathcal{V}'|$. Therefore, there exists an S for which the lemma holds. $\square$

Algorithm

Algorithm 1 shows our full procedure. The main idea is the following: If we pick a committee S that does not lie in $\mathcal{B}_v(\Delta)$ for most voters v, then by Lemma 3.3.4, S_a is forced to lie in $\mathcal{G}_v(\Delta)$ for these voters if S_a beats S. But since Δ is 2-endowment stable, by Lemma 3.3.5, there are only a small number of v where S_a can lie in $\mathcal{G}_v(\Delta)$. We can therefore remove these set of voters for whom $S \notin \mathcal{B}_v(\Delta)$, since S makes sure no S_a can capture too many of these voters. This reduces the number of voters by a constant factor. For the remaining voters, we recursively find another committee of smaller (but not too much smaller) weight, which reduces the number of voters by another constant factor; and so on. The key point is that the total weight of all these committees is a geometric sequence, *and* the number of voters who can be captured by S_a in each round is also a geometric sequence, showing a constant-endowment core solution.

Algorithm 1 Iterated Rounding

1: $t \leftarrow 0$; $\mathcal{V}^{(0)} \leftarrow [n]$; $T^{(0)} \leftarrow \varnothing$; $b^{(0)} \leftarrow (1-\alpha)b$.
2: **while** $\mathcal{V}^{(t)} \neq \varnothing$ **do**
3: $\quad \Delta^{(t)} \leftarrow \textsc{Lottery}(\mathcal{V}^{(t)}, b^{(t)})$.
4: $\quad$ Let $S^{(t)}$ be any committee such that $\left|\left\{v \in \mathcal{V}^{(t)} \mid S^{(t)} \notin \mathcal{B}_v\left(\Delta^{(t)}\right)\right\}\right| \geq (1-\beta) \cdot |\mathcal{V}^{(t)}|$.
5: $\quad \mathcal{W}^{(t)} \leftarrow \left\{v \in \mathcal{V}^{(t)} \mid S^{(t)} \notin \mathcal{B}_v\left(\Delta^{(t)}\right)\right\}$.
6: $\quad \mathcal{V}^{(t+1)} \leftarrow V^{(t)} \setminus \mathcal{W}^{(t)}$.
7: $\quad T^{(t+1)} \leftarrow T^{(t)} \cup S^{(t)}$.
8: $\quad b^{(t+1)} \leftarrow \alpha b^{(t)}$.
9: $\quad t \leftarrow t+1$.
10: **end while**
11: **return** $T^f \leftarrow T^{(t)}$.

Analysis

Line 4 in Algorithm 1 is valid by Lemma 3.3.5. We next bound the weight of the final set.

Lemma 3.3.6. $w(T^f) \leq b$.

Proof. $w(T^f) \leq \sum_{t \geq 1} w(S^{(t)}) \leq \sum_t \alpha^{t-1}(1-\alpha)b \leq b$. □

We finally show that the resulting set is approximately in the core, completing the proof of Theorem 3.1.4.

Lemma 3.3.7. *When $\alpha = \frac{1}{2}$ and $\beta = \frac{1}{4}$, then T^f is a 32-endowment core solution of weight at most b.*

Proof. Given S_a, since $\mathcal{V} = [n] = \bigcup_{t \geq 1} \mathcal{W}^{(t)}$ and $T^f = \bigcup_{t \geq 1} S^{(t)}$, using monotonicity, we have:

$$V\left(T^f, S_a\right) \leq \sum_{t \geq 1} V_{\mathcal{W}^{(t)}}\left(T^f, S_a\right) \leq \sum_{t \geq 1} V_{\mathcal{W}^{(t)}}\left(S^{(t)}, S_a\right).$$

Since $S^{(t)} \notin \mathcal{B}_v\left(\Delta^{(t)}\right)$ for $v \in \mathcal{W}^{(t)}$, by Lemma 3.3.4, $S_a \succ_v S^{(t)}$ only if $S_a \in \mathcal{G}_v\left(\Delta^{(t)}\right)$. By Lemma 3.3.5,

$$\left|\left\{v \in \mathcal{V}^{(t)} \mid S_a \in \mathcal{G}_v\left(\Delta^{(t)}\right)\right\}\right| < \frac{2}{\beta} \cdot \frac{w(S_a)}{b^{(t)}} \cdot |\mathcal{V}^{(t)}|.$$

Note now that $|\mathcal{V}^{(t)}| \leq \beta|\mathcal{V}^{(t-1)}|$, so that $|\mathcal{V}^{(t+1)}| \leq \beta^t n$. Furthermore, $b^{(t)} = \alpha b^{(t-1)}$, so that $b^{(t+1)} = \alpha^t(1-\alpha)b$. Therefore,

$$\begin{aligned} V_{\mathcal{W}^{(t)}}\left(S^{(t)}, S_a\right) &\leq \left|\left\{v \in \mathcal{V}^{(t)} \mid S_a \in \mathcal{G}_v\left(\Delta^{(t)}\right)\right\}\right| \\ &< \frac{2}{\beta} \cdot \frac{w(S_a)}{b^{(t)}} \cdot |\mathcal{V}^{(t)}| \leq \left(\frac{\beta}{\alpha}\right)^{t-1} \cdot \frac{2}{(1-\alpha)\beta} \cdot \frac{w(S_a)}{b} \cdot n. \end{aligned}$$

Summing over all t, we have

$$V\left(T^f, S_a\right) \leq \sum_{t \geq 1} V_{\mathcal{W}^{(t)}}\left(S^{(t)}, S_a\right) < \frac{2\alpha}{\beta(1-\alpha)(\alpha-\beta)} \cdot \frac{w(S_a)}{b} \cdot n.$$

This is minimized when $\beta = \frac{\alpha}{2}$. Setting $\alpha = \frac{1}{2}$, this is at most $32 \cdot \frac{w(S_a)}{b} \cdot n$, completing the proof. □

3.4 67.37-Utility Core for Submodular Utilities

Our overall algorithm proceeds via constructing approximate fractional core for a well-known continuous relaxation, called the *multilinear relaxation*. We use a continuous time local search procedure for the Nash Welfare objective to find this fractional solution. We show that it is almost in a 2-utility fractional core, and subsequently round it iteratively to find a solution in the approximate integer core. Our overall algorithm runs in time polynomial in the number of voters and candidates. We first present the entire algorithm as an existence proof, and then delve into the details of the running time in Section A.1.

3.4.1 First Observations

The following steps are without loss of generality, aiming to simplify future derivations in this section. We will assume $\varepsilon \in (0, 1/20)$ is a small constant that we choose later. In a proof of existence we can set $\varepsilon = 0$; however, we need $\varepsilon > 0$ to achieve polynomial running time.

- Let $u_i^{\max} := \max_j\{u_i(\{j\})\}$, the maximum utility of voter i with a single item. We scale down every $u_i(T)$ by a factor of $u_i^{\max}$ and denote this normalized utility as $u_i'(T) := \frac{u_i(T)}{u_i^{\max}}$. (We ignore every voter with $u_i(C) = 0$, since they will not deviate in any situation.) By submodularity of the utility functions, $u_i'(T) \in [0, m]$ for every i and $T \subseteq C$. This step preserves the original problem, since every voter multiplicatively compares utilities of different committees in Definition 3.1.3 and different voters compare separately. Therefore, without loss of generality, we assume $\max_j\{u_i(\{j\})\} = 1$ and thus $u_i(T) \in [0, m], \forall T \subseteq C$.

- We further assume $s_j \leq b$ for each $j \in C$, as otherwise, the candidate cannot be selected in any feasible committee.

- Now define

$$C_s := \{j \mid s_j \leq \varepsilon b/m\}$$

  as a set of "small enough" candidates. We place all these candidates in the solution. Clearly all candidates in C_s take up a total budget of at most εb – the remaining budget is at least $(1-\varepsilon)b$. Our algorithm then works with candidates in:

$$C_\ell := C \setminus C_s = \{j \mid s_j > \varepsilon b/m\}.$$

  Note that each candidate j in C_ℓ now has size $s_j > \varepsilon b/m$.

3.4.2 The 2-Utility Fractional Core

We now present a procedure that finds an approximate fractional core solution for an arbitrary budget $B \in \left[\frac{\varepsilon}{5m}b, b\right]$, and an arbitrary subset $W \subseteq V$ of voters. We present the algorithm and its optimality properties below; however, we defer the polynomial running time analysis to Section A.1.

Multilinear Extension

Recall that the utility function $u_i(S)$ of voter i for $S \subseteq C$ is non-negative, monotone, and submodular. We first extend these utilities to fractional allocations. A natural way to do that is to use the multilinear extension.

Definition 3.4.1 (Multilinear Extension, [CCPV11, Von08])**.** Given a monotone, non-negative submodular function f, its *multilinear extension* F is defined for any $\mathbf{x} \in [0,1]^m$ as:

$$F(\mathbf{x}) = \sum_{T \subseteq C} f(T) \left(\prod_{j \in T} x_j \right) \left(\prod_{j \notin T} (1 - x_j) \right).$$

We will apply the following property of this function.

Lemma 3.4.2 (Concavity along Positive Directions, [CCPV11])**.** *Given a monotone, non-negative submodular function f, its multilinear extension F is concave along positive directions,* i.e.*, for all $\mathbf{x_0}$ and all $\mathbf{d} \geq \mathbf{0}$, we have that $F_{\mathbf{x_0},\mathbf{d}}(\lambda) = F(\mathbf{x_0} + \lambda \mathbf{d})$ is a monotone, non-negative and concave function of λ.*

We denote the multilinear extension of voter i's utility function u_i for a fractional committee $\mathbf{x}$ as $U_i(\mathbf{x})$. The above lemma implies that $\frac{\partial U_i}{\partial x_j} \geq 0$ for all $i \in W, j \in C$. The next lemma upper bounds the gradient. Here, $\mathbf{x_{-j}}$ is $\mathbf{x}$ with the j^{th} dimension removed, and $T \sim \mathbf{x_{-j}}$ means that T is chosen by including $\ell \in C \setminus \{j\}$ independently with probability x_ℓ.

Lemma 3.4.3 (Bounded Gradient, [CCPV11])**.**

$$\frac{\partial U_i(\mathbf{x})}{\partial x_j} = \sum_{T \subseteq C \setminus \{j\}} \Pr_{T \sim \mathbf{x_{-j}}} [T] \cdot (u_i(T \cup \{j\}) - u_i(T)) \leq \max_j u_i(\{j\}) = 1.$$

The input of the algorithm consists of three parts: The utility function $u_i(T)$ for each committee $T \subseteq C$ and voter $i \in W$, the sizes of the candidates $\{s_j\}$ and the total budget $B \in \left[\frac{\varepsilon}{5m}b, b\right]$ on candidates in C_ℓ. Recall that we assume all candidates in C_s have already been chosen. We assume w.l.o.g. that $\sum_{j\in C_\ell} s_j > B$. In our algorithm, we treat the utility function as an oracle which returns the utility of an agent i on a given committee in $O(1)$ time.

In our algorithm, we will lower bound the allocation x_j by an amount $\underline{x}_j$ that is defined as follows.

$$\underline{x}_j = \begin{cases} 1 & j \in C_s, \\ \frac{B\cdot\varepsilon}{\sum_{j\in C_\ell} s_j} & j \in C_\ell. \end{cases}$$

Nash Welfare. We will find a local optimum to the following *Nash Welfare* program to find the fractional solution.

$$\text{Maximize} \sum_{i\in W} \log U_i(\mathbf{x})$$

$$\begin{array}{rcl} \sum_{j\in C_\ell} s_j x_j & \leq & B \\ x_j & \in & [\underline{x}_j, 1] \quad \forall j \in C \end{array}$$

Let $\phi_i(\mathbf{x}) = \log U_i(\mathbf{x})$, and $\phi(\mathbf{x}) = \sum_{i\in W} \phi_i(\mathbf{x})$. This is the objective value of the above program.

Local Search. We call the local search procedure as $\mathrm{NW}(C, W, \{U_i\}, B)$. This procedure has a *step-size* parameter $\delta = \frac{\varepsilon^7 b}{312m^6}$. We start with any allocation $\mathbf{x_0} \in [\underline{x}_j, 1]^m$ with $\sum_{j\in C_\ell} s_j x_j = B$. Note that such an allocation always exists since we assumed $\sum_{j\in C_\ell} s_j > B$, and further, $\sum_{j\in C_\ell} s_j \underline{x}_j = \varepsilon B$.

Given the current allocation $\mathbf{x}$, we find candidates $j, \ell \in C_\ell$ such that (i) $x_j \leq 1 - \frac{\delta}{s_j}$; (ii) $x_\ell \geq \underline{x}_\ell + \frac{\delta}{s_\ell}$; and (iii) the following condition holds:

$$\frac{\partial \phi(\mathbf{x})}{\partial x_j \cdot s_j} > \frac{\partial \phi(\mathbf{x})}{\partial x_\ell \cdot s_\ell} + \frac{\varepsilon}{b}. \tag{3.5}$$

Recall that b is the budget for the overall problem, not for the subroutine. While such a pair of candidates (j, ℓ) exists, the algorithm increases x_j by δ/s_j, and decreases x_ℓ by δ/s_ℓ. Note that this update is feasible for the program, and preserves the cost of the allocation in C_ℓ at exactly B, and is hence feasible for the above program. The procedure stops when such a pair of candidates (j, ℓ) no longer exists.

Analysis

In the analysis below, we assume the local optimum can be efficiently computed and present properties of this solution. The running time analysis is presented in Section A.1.

We show the following result, which at a high level, extends the analysis of the PAV rule for binary additive utilities and unit sizes in [PS20] to the continuous setting, with submodular utilities and arbitrary sizes. This result will show as a corollary that $\mathrm{NW}(C, W, \{U_i\}, B)$ finds a 2-utility fractional core outcome.

Theorem 3.4.4. *Given a set of voters W with utilities $\{u_i\}$ and multilinear extensions $\{U_i\}$, and a cost budget $B \in \left[\frac{\varepsilon}{5m} b, b\right]$ on C_ℓ, let $\mathbf{x}$ denote the solution to* $\mathrm{NW}(C, W, \{U_i\}, B)$. *Suppose a subset $S \subseteq W$ of voters can choose an allocation $\mathbf{y} \in [0,1]^m$ with $\mathrm{Cost}(\mathbf{y}) \leq b$ such that $U_i(\mathbf{y}) > \theta U_i(\mathbf{x})$ for all $i \in S$, where $\theta \geq 1$. Then:*

$$|S| < \frac{|W|}{B(1-\varepsilon)} \cdot \frac{\mathrm{Cost}(\mathbf{y})}{\theta - 1 - 2\varepsilon}.$$

Proof. Given the local optimum $\mathbf{x}$ with $\sum_{j \in C_\ell} s_j x_j = B$. Let $M_1 = \{j \in C \mid x_j + \frac{\delta}{s_j} \leq$

$y_j\}$ and $M_2 = \{j \in C \mid x_j < y_j < x_j + \frac{\delta}{s_j}\}$. Note that $M_1 \cup M_2 \subseteq C_\ell$, since for all $j \in C_s$, we have $x_j = 1 \geq y_j$.

Let $\mathbf{y}'$ be the allocation where $y'_j = y_j$ for $j \in M_1$, and $y'_j = x_j$ otherwise. Note that when moving from $\mathbf{x}$ to $\mathbf{y}'$, utility increase for $i \in S$ can only come from allocation increase of candidates in M_1. Further, note that the direction $\mathbf{y}' - \mathbf{x}$ is non-negative, since $y_j > x_j$ for $j \in M_1$ and $y'_j = x_j$ for $j \notin M_1$.

By applying Lemma 3.4.2 to the multilinear function U_i and observing that $\frac{\partial U_i}{\partial x_j} \geq 0$, we have

$$\sum_{j \in M_1} \frac{\partial U_i(\mathbf{x})}{\partial x_j} \cdot y_j + U_i(\mathbf{x}) \geq \sum_{j \in C} \frac{\partial U_i(\mathbf{x})}{\partial x_j} \cdot (y'_j - x_j) + U_i(\mathbf{x}) \geq U_i(\mathbf{y}'),$$

where the last inequality follows from the concavity of U_i in non-negative directions.

Consider now the allocation $\mathbf{z}$, where $z_j = y_j$ if $j \in M_1 \cup M_2$, and $z_j = x_j$ otherwise. Note that $\mathbf{z} \geq \mathbf{y}$, so that $U_i(\mathbf{z}) \geq U_i(\mathbf{y})$. Further, the direction $\mathbf{z} - \mathbf{y}'$ is non-negative, since $z_j = y_j \geq x_j = y'_j$ for $j \in M_2$, and $z_j = y'_j$ otherwise. Applying Lemma 3.4.2 again gives

$$U_i(\mathbf{z}) \leq U_i(\mathbf{y}') + \sum_j \frac{\partial U_i(\mathbf{y}')}{\partial y'_j} \cdot (z_j - y'_j) = U_i(\mathbf{y}') + \sum_{j \in M_2} \frac{\partial U_i(\mathbf{y}')}{\partial x_j} \cdot (y_j - y'_j).$$

Combining the relations above and observing that $U_i(\mathbf{z}) \geq U_i(\mathbf{y})$, we have

$$\sum_{j \in M_1} \frac{\partial U_i(\mathbf{x})}{\partial x_j} \cdot y_j + U_i(\mathbf{x}) \geq U_i(\mathbf{y}') \geq U_i(\mathbf{y}) - \sum_{j \in M_2} \frac{\partial U_i(\mathbf{y}')}{\partial x_j} \cdot (y_j - y'_j).$$

By Lemma 3.4.3, we have $\frac{\partial U_i(\mathbf{y}')}{\partial x_j} \leq 1$ for all $j \in M_2$. Further, $y_j - y'_j = y_j - x_j \leq \frac{\delta}{s_j} \leq \frac{\varepsilon^6}{312m^5}$, since $s_j \geq \frac{\varepsilon}{m} b$ and $\delta = \frac{\varepsilon^7 b}{312m^6}$. Therefore, for any $i \in S$, we have:

$$\sum_{j \in M_1} \frac{\partial U_i(\mathbf{x})}{\partial x_j} \cdot y_j \geq U_i(\mathbf{y}) - U_i(\mathbf{x}) - \frac{\varepsilon^6}{312m^4} > (\theta - 1)U_i(\mathbf{x}) - \frac{\varepsilon^6}{312m^4},$$

where we have used that $U_i(\mathbf{y}) > \theta U_i(\mathbf{x})$ for all $i \in S$.

Note now that

$$U_i(\mathbf{x}) \geq u_i^{\max} \cdot \underline{x}_j = \frac{B \cdot \varepsilon}{\sum_{j \in C_\ell} s_j} \geq \frac{\varepsilon^2 \cdot b/m}{b \cdot m} = \frac{\varepsilon^2}{m^2}, \tag{3.6}$$

where we have used that $s_j \leq b$ for all $j \in C$, and $B \geq \frac{\varepsilon}{5m} b$. Combining the two inequalities above gives

$$\sum_{j \in M_1} \frac{\partial U_i(\mathbf{x})}{\partial x_j} \cdot y_j \geq (\theta - 1 - \varepsilon) \cdot U_i(\mathbf{x}).$$

Summing this inequality over all candidates in S, we have

$$\sum_{i \in S} \sum_{j \in M_1} \frac{\partial U_i(\mathbf{x})}{\partial x_j} \cdot \frac{y_j}{U_i(\mathbf{x})} \geq |S| \cdot (\theta - 1 - \varepsilon).$$

Since $\frac{\partial \phi_i(\mathbf{x})}{\partial x_j} = \frac{1}{U_i(\mathbf{x})} \cdot \frac{\partial U_i(\mathbf{x})}{\partial x_j}$, we have

$$\begin{aligned} \frac{\sum_{j \in M_1} \left(y_j \cdot s_j \cdot \frac{\partial \phi(\mathbf{x})}{\partial x_j \cdot s_j} \right)}{\sum_{j \in M_1} y_j \cdot s_j} \geq \frac{\sum_{j \in M_1} \left(\sum_{i \in S} \frac{\partial \phi_i(\mathbf{x})}{\partial x_j \cdot s_j} \cdot y_j \cdot s_j \right)}{\sum_{j \in M_1} y_j \cdot s_j} &\geq \frac{\sum_{i \in S} \sum_{j \in M_1} \frac{\partial U_i(\mathbf{x})}{\partial x_j} \cdot \frac{y_j}{U_i(\mathbf{x})}}{\sum_{j \in M_1} y_j \cdot s_j} \\ &\geq \frac{|S| \cdot (\theta - 1 - \varepsilon)}{\mathrm{Cost}(\mathbf{y})}. \end{aligned}$$

Therefore,

$$\max_{j \in M_1} \frac{\partial \phi(\mathbf{x})}{\partial x_j \cdot s_j} > \frac{|S| \cdot (\theta - 1 - \varepsilon)}{\mathrm{Cost}(\mathbf{y})} \geq \frac{|S| \cdot (\theta - 1 - 2\varepsilon)}{\mathrm{Cost}(\mathbf{y})} + \varepsilon/b,$$

where we have used that $\mathrm{Cost}(\mathbf{y}) \leq b$ and $|S| \geq 1$. Let j_1 denote the $j \in M_1$ that achieves the maximum in the previous inequality.

Now, let $R_1 = \{j \in C_\ell \mid x_j \geq \underline{x}_j + \frac{\delta}{s_j}\}$. Applying Lemma 3.4.2 along the positive direction $\mathbf{x}$ gives $\sum_{j \in C} x_j \cdot \frac{\partial U_i(\mathbf{x})}{\partial x_j} \leq U_i(\mathbf{x})$. Since $\frac{\partial U_i(\mathbf{x})}{\partial x_j} \geq 0$ for all $j \in C$, we have

$$\sum_{j \in R_1} x_j \cdot \frac{\partial U_i(\mathbf{x})}{\partial x_j} \leq U_i(\mathbf{x}).$$

Summing this over all $i \in W$, we have:

$$\sum_{i\in W}\sum_{j\in R_1} x_j \cdot s_j \cdot \frac{\partial \phi_i(\mathbf{x})}{\partial x_j \cdot s_j} = \sum_{i\in W} \frac{\sum_{j\in R_1} x_j \cdot \frac{\partial U_i(\mathbf{x})}{\partial x_j}}{U_i(\mathbf{x})} \leq \sum_{i\in W} \frac{U_i(\mathbf{x})}{U_i(\mathbf{x})} = |W|.$$

Therefore,

$$\frac{\sum_{j\in R_1} x_j \cdot s_j \cdot \frac{\partial \phi(\mathbf{x})}{\partial x_j \cdot s_j}}{\sum_{j\in R_1} x_j \cdot s_j} = \frac{\sum_{i\in W}\sum_{j\in R_1} x_j \cdot s_j \cdot \frac{\partial \phi_i(\mathbf{x})}{\partial x_j \cdot s_j}}{\sum_{j\in R_1} x_j \cdot s_j} \leq \frac{|W|}{\sum_{j\in R_1} x_j \cdot s_j} \leq \frac{|W|}{B(1-\varepsilon)}.$$

The last inequality holds since $\sum_{j\in C_\ell \setminus R_1} x_j \cdot s_j \leq \sum_{j\in C_\ell \setminus R_1} \frac{\delta}{s_j} \cdot s_j \leq m\delta \leq \varepsilon B$, where we have used that $\delta = \frac{\varepsilon^7 b}{312 m^6}$ and $B \geq \frac{\varepsilon}{5m} b$. This implies $\sum_{j\in R_1} s_j x_j \geq B(1-\varepsilon)$.

Therefore,

$$\min_{j\in R_1} \frac{\partial \phi(\mathbf{x})}{\partial x_j \cdot s_j} \leq \frac{|W|}{B(1-\varepsilon)}.$$

Let j_2 denote the $j \in R_1$ that achieves this minimum.

Note that $x_{j_1} < y_{j_1}$ and $x_{j_2} > \underline{x}_{j_2}$ by assumption. Further, since $\mathbf{x}$ is locally optimal, Eq. (3.5) cannot hold for $j = j_2$ and $\ell = j_1$. This means $\frac{\partial \phi(\mathbf{x})}{\partial x_{j_1} \cdot s_{j_1}} \leq \frac{\partial \phi(\mathbf{x})}{\partial x_{j_2} \cdot s_{j_2}} + \varepsilon/b$. Therefore, we have

$$\frac{|S| \cdot (\theta - 1 - 2\varepsilon)}{\mathrm{Cost}(\mathbf{y})} < \frac{\partial \phi(\mathbf{x})}{\partial x_{j_1} \cdot s_{j_1}} - \varepsilon/b \leq \frac{\partial \phi(\mathbf{x})}{\partial x_{j_2} \cdot s_{j_2}} \leq \frac{|W|}{B(1-\varepsilon)}.$$

Rearranging completes the proof. □

As an aside, the following corollary is immediate if all one wants is an approximate fractional core on the utilities $\{U_i\}$ and budget b. Note that the proof trivially extends to *any* continuous, concave utilities $\{U_i\}$, since these are concave in positive directions.

Corollary 3.4.5. *The solution* $\mathbf{x}$ *of* $\mathrm{NW}(C, V, \{U_i\}, b)$ *is in a* 2*-utility fractional core for* any *continuous concave utilities* $\{U_i\}$.

Proof. Setting $\theta = 2$ and $\varepsilon = 0$ in Theorem 3.4.4, if there exists an S and a solution $\mathbf{y}$ such that $U_i(\mathbf{y}) > 2U_i(\mathbf{x})$ for all $i \in S$, then $|S| < \frac{|W|}{b} \cdot \text{Cost}(\mathbf{y})$. This contradicts Definition 3.2.2 and thus completes the proof. $\square$

3.4.3 Randomized Rounding and Satisfied Voters

We now present a randomized rounding scheme on the fractional solution constructed above. Since the utilities of voters are different but the allocation is the same, we will need the rounding of the fractional allocation to be *oblivious* to the utilities. The natural approach is therefore randomized rounding. Clearly, if we randomly round, the expected utility of a voter in the integral committee is at least that in the fractional core. However, the same cannot be said for an arbitrary *realization* of the randomization: there can be many voters having lower-than-expected utilities, and they can form a coalition and deviate. Further, the size of the resulting committee can exceed the bound B. Therefore, we cannot directly argue that the resulting integral committee will lie in the core.

Our contribution in this section is to show a rounding procedure that ensures that with constant probability, a constant fraction of voters obtain at least a constant fraction of the utilities they obtained in the fractional solution. We will use this as a subroutine in our main procedure in Section 3.4.4 in order to argue that these voters will not form part of any deviating coalition.

Rounding Procedure

To begin with, in the sequel, we will distinguish between items that are fractionally and integrally allocated as follows:

Definition 3.4.6 (Fully and Fractionally Allocated Items)**.** An item $j \in C$ in the fractional solution is *fully allocated* if $x_j \geq 1$; we call it *fractionally allocated* if $x_j \in (0, 1)$. Note that all items in C_s are fully allocated.

The rounding procedure has a parameter $\kappa \leq 1$, and executes as shown in Algorithm 2.

Algorithm 2 Randomized Rounding of Fractional Allocation $\mathbf{x}$

```
function Round(C, W, {U_i}, B)
    x ← NW(C, W, {U_i}, κB)
    O ← C_s
    T_2 ← {j ∈ C_ℓ | s_j ≤ κB}
    Include j ∈ T_2 in O independently with probability min(1, x_j)
    return O
end function
```

Satisfied Voters

We now define what it means for a voter to be *satisfied*, and present a bound on the number of satisfied voters after the procedure Round is applied.

Definition 3.4.7 (γ-Satisfied)**.** Let $\mathbf{x}$ be the fractional solution of $\mathrm{NW}(C, W, \{U_i\}, \kappa B)$, and O be the resulting integer committee found by $\textsc{Round}(C, W, \{U_i\}, B)$. For $\gamma \geq 2$, we say that voter i is *γ-satisfied* by O w.r.t. solution $\mathbf{x}$ if there is a $q \in C$ such that $u_i(O \cup \{q\}) \geq \frac{U_i(\mathbf{x})}{\gamma}$.

Our main result in this section is the following theorem.

Theorem 3.4.8 (Constant Fraction of Constant-Satisfied Voters for Submodular Utilities)**.** *Given the fractional solution* $\mathbf{x}$ *produced by* $\mathrm{NW}(C, W, \{U_i\}, \kappa B)$ *where* $|W| = n'$*, there is a integral committee* O *produced by* $\textsc{Round}(C, W, \{U_i\}, B)$ *with at least* $(1-\beta)n'$ *γ-satisfied voters, where*

$$\beta = \left(\kappa e^{1-\kappa}\right)^{\frac{1}{\kappa}} + (\gamma - 1)e^{2-\gamma}. \tag{3.7}$$

Further, this committee is feasible, so that $\sum_{j \in O \cap C_\ell} s_j \leq B$.

For the proof, we will need the following lemma, which is the analog of Lemma A.2.7 for submodular functions.

Lemma 3.4.9 (Lower Tail Bound, [CVZ10]). *Assume $f(\{j\}) \in [0,1]$ for all $j \in C$, and let $\mu_0 = F(\mathbf{x})$, where F is the multilinear extension of f. Let O denote the outcome of independent randomized rounding applied to $\mathbf{x} \in [0,1]^m$. Then for any $\rho > 1$:*

$$\Pr\left[f(O) \leq \frac{\mu_0}{\rho}\right] \leq \left(\rho e^{1-\rho}\right)^{\frac{\mu_0}{\rho}}.$$

Proof of Theorem 3.4.8. We show that for each voter i, the probability that i is γ-satisfied is at least $1-\beta$. This implies the number of γ-satisfied voters is at least $(1-\beta)n'$ in expectation, and thus it is at least $(1-\beta)n'$ for some realization produced by Round.

Let $T_1 = \{j \in C_\ell \mid s_j \geq \kappa B\}$ and $T_2 = C \setminus T_1$. Fix voter i with utility function u_i, and multilinear extension U_i.

Let G_1 denote the multilinear extension restricted only to candidates in T_1, that is,

$$G_1(\mathbf{x}) = \sum_{T \subseteq T_1} u_i(T) \left(\prod_{j \in T} x_j\right) \left(\prod_{j \in T_1 \setminus T} (1-x_j)\right),$$

and similarly, G_2 denote the multilinear extension restricted to T_2, that is,

$$G_2(\mathbf{x}) = \sum_{T \subseteq T_2} u_i(T) \left(\prod_{j \in T} x_j\right) \left(\prod_{j \in T_2 \setminus T} (1-x_j)\right).$$

By sub-additivity of u_i, we have $U_i(\mathbf{x}) \leq G_1(\mathbf{x}) + G_2(\mathbf{x})$, and by monotonicity of u_i, we have $U_i(\mathbf{x}) \geq \max\left(G_1(\mathbf{x}), G_2(\mathbf{x})\right)$.

We now split the analysis into several cases:

Case (1): Suppose $G_1(\mathbf{x}) \geq \frac{1}{\gamma} U_i(\mathbf{x})$. In this case, we make two observations. First,

$$G_1(\mathbf{x}) \leq \sum_{T \subseteq T_1} \left(\sum_{j \in T} u_i(\{j\})\right) \left(\prod_{j \in T} x_j\right) \left(\prod_{j \in T_1 \setminus T} (1-x_j)\right) \leq \sum_{j \in T_1} x_j u_i(\{j\}).$$

Further, since $\sum_{j \in T_1} s_j x_j \leq \kappa B$, and since $s_j \geq \kappa B$ for all $j \in T_1$, we have $\sum_{j \in T_1} x_j \leq 1$. Putting these together, we have

$$G_1(\mathbf{x}) \leq \max_{j \in T_1} u_i(\{j\}).$$

Therefore, using $j^* = \operatorname{argmax}_{j \in T_1} u_i(\{j\})$ as additament makes i achieve utility at least $1/\gamma$ fraction of $U_i(\mathbf{x})$. Therefore, i is γ-satisfied just by the additament.

Case (2): We now assume $G_1(\mathbf{x}) \leq \frac{1}{\gamma} U_i(\mathbf{x})$ and thus $G_2(\mathbf{x}) \geq (1 - 1/\gamma) U_i(\mathbf{x})$. Suppose there exists $\ell \in T_2$ such that $u_i(\{\ell\}) \geq \frac{1}{\gamma - 1} G_2(\mathbf{x})$. Then, just using ℓ as additament causes the utility of i to be at least $1/\gamma$ fraction of $F_i(\mathbf{x})$ and voter i is γ-satisfied by the additament ℓ.

Case (3): In the final case, we have $G_2(\mathbf{x}) \geq (1 - 1/\gamma) U_i(\mathbf{x})$, and we assume that for all $j \in T_2$, $u_i(\{\ell\}) < \frac{1}{\gamma - 1} G_2(\mathbf{x})$. Let $O' = O \cap C_\ell$.

First note that $\mathbb{E}\left[\sum_{j \in O'} s_j\right] \leq \kappa B$, since O' is the result of randomized rounding of $\mathbf{x}$ and since $\sum_{j \in C_\ell} s_j x_j \leq \kappa B$. Further, $O' \subseteq T_2 \cap C_\ell$, and for all $j \in T_2 \cap C_\ell$, we have $s_j \leq \kappa B$. Therefore, a standard application of Chernoff bounds yields:

$$\Pr\left[\sum_{j \in O'} s_j > B\right] \leq \left(\kappa e^{1-\kappa}\right)^{\frac{1}{\kappa}}.$$

Further, note that for all $j \in T_2$, we have that the marginal $u_i(\{\ell\}) < \frac{1}{\gamma - 1} G_2(\mathbf{x})$. Applying Lemma 3.4.9 with $\rho = \gamma - 1$, we have:

$$\Pr\left[u_i(O) \leq \frac{G_2(\mathbf{x})}{\gamma - 1}\right] \leq (\gamma - 1) e^{2-\gamma}.$$

Therefore, by union bounds, with probability at least $1 - \beta$, we have both events: (1) The committee O' is feasible for size B; and (2) $u_i(O) \geq \frac{G_2(\mathbf{x})}{\gamma - 1} \geq \frac{U_i(\mathbf{x})}{\gamma}$, implying i is γ-satisfied by O. This completes the proof. $\square$

The following corollary shows that this step can be implemented in polynomial time.

Corollary 3.4.10. *For β defined in Eq.* (3.7)*, $n' = |W|$ and any $\varepsilon \in (0,1)$, a committee $O \subseteq C$ with $(1-\beta-\varepsilon)n'$ γ-satisfied voters can be computed with probability $1 - \frac{1}{poly(m,n)}$ in time $poly(n,m,1/\varepsilon)$.*

Proof. Let X be the random variable indicating the number of γ-satisfied voters returned by $\textsc{Round}(C, W, \{U_i\}, B)$. By Theorem 3.4.8, we have $\mathbb{E}[X] = (1-\beta)n'$. Further, $X \leq n'$. Therefore, $\Pr[X \geq (1-\beta-\varepsilon)n] \geq \frac{\varepsilon}{\beta+\varepsilon} \geq \frac{\varepsilon}{2}$. A standard application of sampling bounds now completes the proof. $\square$

3.4.4 The Constant-Utility Core

In this section, we design an algorithm for the overall problem, and prove that it returns a committee in an $O(1)$-utility core. The algorithm will repeatedly construct fractional solutions and round them, using the algorithms developed above as subroutines. The analysis critically requires the local optimality property of the Nash Welfare objective, captured in Theorem 3.4.4.

Algorithm

Algorithm 3 finds an approximate core solution in the following way: It iteratively computes the fractional local optimum to Nash Welfare on the remaining voters with a scaled-down budget, rounds it, eliminates voters that are γ-satisfied with respect to the solution at this iteration, scales down the budget again, and iterates on the remaining voters with the smaller budget. The scaling parameter ω for the budget will be determined later. The overall structure of the algorithm is similar to that in [JMW20], though the details of constructing the fractional solution, and the resulting proof of correctness are entirely different.

Algorithm 3 Iterative Rounding of the Nash Welfare Solution

function $\textsc{IterRound}(C, V, \{u_{ij}\}, b)$
  $t \leftarrow 0$
  $V_0 \leftarrow V$
  $T^* \leftarrow C_s$
  $b_0 \leftarrow (1-\varepsilon)(1-\omega)b$
  while $b_t \geq \frac{\varepsilon}{m} b$ **do**
    $\mathbf{x}_t \leftarrow \mathrm{NW}(C, V_t, \{U_i\}, \kappa b_t)$
    $O_t \leftarrow$ Solution of $\textsc{Round}(C, W, \{U_i\}, b_t)$ that satisfies Corollary 3.4.10
    $W_t \leftarrow$ Voters in V_t that are γ-satisfied by O_t with respect to the solution $\mathbf{x}_t$
    $V_{t+1} \leftarrow V_t \setminus W_t$
    $T^* \leftarrow T^* \cup O_t$
    $b_{t+1} \leftarrow \omega b_t$
    $t \leftarrow t+1$
  end while
  return T^* as the final integer solution
end function

Analysis

Let $n = |V|$, and recall that b is the initial budget. We will show that the solution T^* lies in the α-utility core for the set V of voters with size constraint b. First note that since $\sum_{j \in C_s} s_j \leq \varepsilon b$, and since each O_t is feasible for budget b_t, we have:

$$\sum_{j \in T^*} s_j \leq \varepsilon b + \sum_t \sum_{j \in O_t} s_j \leq \varepsilon b + \sum_t b_t = \varepsilon b + (1-\varepsilon)(1-\omega) \sum_{t \geq 0} \omega^t b \leq b.$$

Therefore, the solution T^* is feasible for the size b.

First note that when $b_t < \frac{\varepsilon}{m} b$, since all items in C_ℓ has $s_j > \frac{\varepsilon}{m} b$, we have $\sum_{j \in C_\ell} x_{jt} < 1$. Since $C_s \subseteq T^*$, for any fractional solution $\mathbf{x_t}$, we have $\max_{j \in C} u(\{j\}) \geq U_i(\mathbf{x_t})$, so that all voters are 1-satisfied. This implies $V_t = \varnothing$ at termination, so that any voter $i \in V$ belongs to $W_{t'}$ for some t'.

For the purpose of contradiction, we assume the resulting solution T^* is not in the α-utility core (Definition 3.1.3), where $\alpha \geq 1$ is a quantity we will determine later. Let S denote the set of voters that deviate, and let A denote the set of items they deviate to. We have $\sum_{j \in A} s_j \leq \frac{|S|}{n} \cdot b$, and A provides an α-factor larger utility to voters in S even after including any additament.

Consider the voters in W_t, and let $S_t = S \cap W_t$. These voters are γ-satisfied by O_t with respect to the fractional solution $\mathbf{x}_t$ (Definition 3.4.7). Therefore, if $i \in S_t$ deviates to A to obtain an α-factor larger utility, it must be that $U_i(A) \geq \frac{\alpha}{\gamma} \cdot U_i(\mathbf{x}_t)$. Let $\theta = \frac{\alpha}{\gamma}$. We will assume $\theta > 1$ below.

Using Theorem 3.4.4, since $\mathbf{x}_t$ is a local optimum, using $\mathbf{y}$ as the set S_t, and observing that all agents $i \in S_t$ have $U_i(\mathbf{y}) > \theta U_i(\mathbf{x}_t)$, where $\theta = \frac{\alpha}{\gamma}$, we have:

$$|S_t| \leq \frac{n_t}{\kappa b_t \cdot (1-\varepsilon)} \cdot \frac{\sum_{j \in A} s_j}{\theta - 1 - 2\varepsilon}.$$

Summing this over all t, and using $\sum_{j \in A} s_j \leq \frac{|S|}{n} b$, we have:

$$|S| = \sum_t |S_t| \leq \frac{\sum_{j \in A} s_j}{(\theta - 1 - 2\varepsilon)(1-\varepsilon)} \cdot \sum_t \frac{n_t}{\kappa b_t} \leq \frac{|S|}{n} \cdot \frac{b}{(\theta - 1 - 2\varepsilon)(1-\varepsilon)} \cdot \sum_t \frac{n_t}{\kappa b_t}.$$

Therefore, for a blocking coalition to exist, we need:

$$\frac{b}{n} \cdot \sum_t \frac{n_t}{\kappa b_t} \geq (\theta - 1 - 2\varepsilon)(1-\varepsilon) = \left(\frac{\alpha}{\gamma} - 1 - 2\varepsilon\right)(1-\varepsilon). \tag{3.8}$$

We will now set the parameters ω, γ, α so that the above inequality is false. First note by Corollary 3.4.10 that $n_{t+1} \leq (\beta + \varepsilon) n_t$ where β satisfied Eq. (3.7). Further, $b_{t+1} = \omega b_t$. Therefore,

$$\frac{n_{t+1}}{b_{t+1}} \leq \frac{\beta + \varepsilon}{\omega} \cdot \frac{n_t}{b_t}$$

with $\frac{n_0}{b_0} = \frac{n}{(1-\omega)(1-\varepsilon)b}$. Therefore,

$$\frac{(1-\varepsilon)b}{n} \cdot \sum_t \frac{n_t}{\kappa b_t} \leq \frac{1}{(1-\omega)\kappa} \cdot \sum_{t \geq 0} \left(\frac{\beta + \varepsilon}{\omega}\right)^t = \frac{\omega}{(1-\omega)(\omega - \beta - \varepsilon)\kappa}. \tag{3.9}$$

Combining Eq. (3.8) and Eq. (3.9), for a blocking coalition to exist, we need

$$\alpha \leq \frac{\omega\gamma}{\kappa(1-\omega)\left(\omega-(\gamma-1)e^{2-\gamma}-(\kappa e^{1-\kappa})^{\frac{1}{\kappa}}-\varepsilon\right)(1-\varepsilon)^2}+(1+2\varepsilon)\cdot\gamma.$$

For an α slightly larger than the right-hand side, a blocking coalition will therefore not exist. Then The right-hand side of the above inequality is approximately minimized when $\omega = 0.23$, $\gamma = 7.435$, $\kappa = 0.21$ and $\varepsilon \to 0$, yielding $\alpha < 67.37$.[1] This finally yields the following theorem; the only missing detail is the running time of local search in Section 3.4.2, which we address in Section A.1.

Theorem 3.1.5 (Main Theorem for Utility Approximation)**.** *For monotone submodular utilities, a* 67.37*-utility core is always non-empty. One such solution can be computed in polynomial time.*

3.5 Lower Bounds on Approximations

In this section, we provide lower bound examples, complementing our main theorems of Theorem 3.1.4 and Theorem 3.1.5.

3.5.1 Lower Bound for α-Endowment Core

It is relatively easy to construct a *participatory budgeting* preference profile where a $(2-\varepsilon)$-endowment core solution does not exist. This instance has cyclic preferences. There are m candidates $\{c_i\}_{i\in[m]}$ of unit weight, and $n = m$ voters $\{v_i\}_{i\in[n]}$. Let $b = 2 - \frac{\varepsilon}{2}$. The preference of v_i is:

$$c_i \succ c_{i+1} \succ \cdots \succ c_m \succ c_1 \succ \cdots \succ c_{i-1}.$$

Any feasible committee is some single candidate c_i, but all voters except v_i can deviate and choose c_{i-1} (or c_m if $i = 1$). Therefore, the approximation ratio is at least $\frac{m-1}{n} \cdot b = \frac{m-1}{m} \cdot \left(2 - \frac{\varepsilon}{2}\right) > 2 - \varepsilon$ when m is large enough.

[1] Note that for this choice of κ, the subroutine $\mathrm{NW}(C, V_t, \{U_i\}, \kappa b_t)$ is run with budget at least $\kappa b_t \geq \kappa\frac{\varepsilon}{m}b \geq \frac{\varepsilon}{5m}b$, so that the precondition of Theorem 3.4.4 holds.

We now strengthen this example to show that even for the RANKING or additive-utility setting with unit-weight candidates and integral committee size limit b, there exist instances where a $(2-\varepsilon)$-endowment core is empty. In the RANKING setting, a voter prefers a committee S_1 to S_2 if and only if she prefers some $c_1 \in S_1$ to any $c_2 \in S_2$. In other words, when evaluating any committee, a voter only cares about her favorite candidate in that committee.

Theorem 3.1.7 (Endowment Core Lower Bound)**.** *In the unweighted* RANKING *or additive-utility setting,* $(2-\varepsilon)$*-endowment core of integral size limit* b *may not exist for any constant* $\varepsilon > 0$*.*

Proof. We first prove the theorem for the RANKING setting. For any positive integer $r \geq 2$ and ℓ, we construct the following instance with $n = r \cdot \ell$ voters and $m = r \cdot \ell$ candidates. We view the candidates $\{c_{i,j}\}_{i\in[r],j\in[\ell]}$ as a matrix with r rows and ℓ columns. The voters are $\{v_{i,j}\}_{i\in[r],j\in[\ell]}$, where $v_{i,j}$ has the following preference:

- For candidates not in the same row, her preference is

$$c_{i,j_i} \succ c_{i+1,j_{i+1}} \succ \cdots \succ c_{r,j_r} \succ c_{1,j_1} \succ \cdots \succ c_{i-1,j_{i-1}}$$

  for any $j_1, j_2, \cdots, j_r$.

- For candidates in the same row i', her preference is

$$c_{i',j} \succ c_{i',j+1} \succ \cdots \succ c_{i',\ell} \succ c_{i',1} \succ \cdots \succ c_{i',j-1}.$$

Let $b = r - 1$. For any deterministic committee S_d of size b, there must be some $i \in [r]$, so that no candidate from the i-th row is in S_d, and at most one candidate from the $(i+1)$-th (or first if $i = r$) row is in S_d. Otherwise, every row where no candidate is selected must be followed by a row where at least 2 candidate is selected. Thus, on average, at least 1 candidate is selected from each row, contradicting with $b = r - 1$.

Let the candidate in the $(i+1)$-th row in S_d be $c_{i+1,j}$ if there is one. Notice that $c_{i+1,j-1}$ is preferred to S_d by at least $2\ell - 1$ voters: those in the i-th row and the $(i+1)$-th row except $v_{i+1,j}$. Therefore, the approximation ratio c is at least $\frac{(2\ell-1)b}{n} = \frac{(2\ell-1)(r-1)}{\ell r}$, which is close to 2 when ℓ and r are large.

For the additive-utility setting, we use the same instance, except that every voter assigns utility of $2^m, 2^{m-1}, \ldots, 2$ for the candidates according to her preference ordering. The new preference on committees is strictly more refined than the one in the Ranking setting, and therefore the instance gives a lower bound for the additive-utility setting too. □

3.5.2 *Lower Bound for α-Utility Core: General Monotone Utilities*

We restate our lower bound for general monotone utilities for the α-utility core. This rules out extending our constant factor bound to general monotone utilities.

Theorem 3.1.8 (General Utility Lower Bound)**.** *For general monotone utilities, for any function $\varphi : \mathbb{Z}^+ \times \mathbb{Z}^+ \to \mathbb{R}^+$, a $\varphi(n,m)$-utility core can be empty.*

Theorem 3.1.8 is proved by the following example. (The same structure of 2 groups of 3 cyclically symmetric voters appears in [FMS18, PPS21].)

Example 3.5.1. We have $n = 6$ voters and $m = 30$ candidates. The candidates are grouped into 6 disjoint sets, each of which contains 5 candidates and is called a "gadget". We name the gadgets $g_1, \ldots, g_6$. Each voter i has a favorite gadget g_{f_i} and a second favorite gadget g_{s_i}, given by:

$$f_1 = 1,\ f_2 = 2,\ f_3 = 3,\ f_4 = 4,\ f_5 = 5,\ f_6 = 6;$$

$$s_1 = 2,\ s_2 = 3,\ s_3 = 1,\ s_4 = 5,\ s_5 = 6,\ s_6 = 4.$$

For any committee E, let $x_i(E) = \frac{1}{5}|E \cap g_{f_i}|$ and $y_i(E) = \frac{1}{5}|E \cap g_{s_i}|$, denoting the fraction of candidates in the favorite / second favorite gadget of voter i being selected into E, respectively. The utility of voter i on E is given by

$$u_i(E) = (\alpha + 1) \cdot \mathbb{1}[x_i(E) = 1] + \mathbb{1}[y_i(E) = 1].$$

Here $\alpha = \varphi(n, m) = \varphi(6, 30)$. Her utility is monotone and supermodular.

Each candidate is of unit size 1 and the budget $b = 15$. For any feasible committee E, there must be at least 3 gadgets g's with $|E \cap g| \leq 3$ – otherwise the committee has at least $4 \cdot 4 = 16 > b$ candidates. Therefore, either $\{g_1, g_2, g_3\}$ or $\{g_4, g_5, g_6\}$ includes at least 2 gadgets with $|E \cap g| \leq 3$. Without loss of generality, assume $\{g_1, g_2, g_3\}$ does and $|E \cap g_1| \leq 3$, $|E \cap g_2| \leq 3$. In this case, voters 1 and 2 can deviate and buy g_2, as they have a budget of $b \cdot \frac{2}{n} = 5$. For any additaments q and q':

$$u_1(g_2) = 1, \qquad u_1(E \cup \{q\}) = 0;$$
$$u_2(g_2) = \alpha + 1, \quad u_2(E \cup \{q'\}) \leq 1.$$

We have $u_1(g_2) > \alpha u_1(E \cup \{q\})$ and $u_2(g_2) > \alpha u_2(E \cup \{q'\})$.

3.5.3 Lower Bound for α-Utility Core: Submodular Utilities

Next, we modify Example 3.5.1 to show a lower bound for monotone submodular utilities. We show that for submodular utilities, there is a lower bound on approximation of an absolute constant $\alpha > 1$, and therefore the multiplicative factor is necessary for the α-utility core to always exist.

Theorem 3.1.9 (Submodular Lower Bound)**.** *For monotone submodular utilities, a* 1.015*-utility core can be empty.*

Example 3.5.2. We use the same setting as Example 3.5.1, except the utility of each voter i is given by

$$u_i(E) = x_i(E) + z \cdot (1 - x_i(E)) \cdot y_i(E),$$

where $z \in (0, 1)$ is a constant to be determined later.

Lemma 3.5.3. *The function u_i is monotone and submodular.*

Proof. Fix any $E \subseteq C$ and $t \notin E$ and consider $u_i(E \cup \{t\}) - u_i(T)$. Since $g_{f_i} \cap g_{s_i} = \varnothing$, this t lies in one of these two sets but not both. Suppose $t \in g_{s_i}$. Then,

$$u_i(E \cup \{t\}) - u_i(E) = z \cdot (1 - x_i(E)) \cdot (y_i(E \cup \{t\}) - y_i(E)) \geq 0.$$

Similarly, if $t \in g_{f_i}$, we have

$$u_i(E \cup \{t\}) - u_i(E) = (1 - zy_i(E)) \cdot (x_i(E \cup \{t\}) - x_i(E)) \geq 0,$$

where we have used that $y_i(E) \leq 1$ and $z \in [0, 1]$. This shows that u_i is a monotone function.

Similarly, if $E \subseteq E'$, then $1 - x_i(E) \geq 1 - x_i(E')$ since the coverage function is monotone. Further, by the submodularity of the coverage function, we have

$$y_i(E \cup \{t\}) - y_i(E) \geq y_i(E' \cup \{t\}) - y_i(E').$$

Therefore, if $t \in g_{s_i}$, then

$$\begin{aligned} u_i(E \cup \{t\}) - u_i(E) &= z \cdot (1 - x_i(E)) \cdot (y_i(E \cup \{t\}) - y_i(E)) \\ &\geq z \cdot (1 - x_i(E')) \cdot (y_i(E' \cup \{t\}) - y_i(E')) \\ &= u_i(E' \cup \{t\}) - u_i(E'). \end{aligned}$$

A similar argument for the case where $t \in g_{f_i}$ completes the proof of submodularity.

□

For any feasible committee E, again without loss of generality, assume $|E \cap g_1| \leq 3$, $|E \cap g_2| \leq 3$. In this case, voters 1 and 2 can deviate and buy g_2. For any additaments q and q':

$$u_1(g_2) = z, \quad u_1(E \cup \{q\}) \leq 0.8 + z \cdot 0.2 \cdot 0.6;$$

$$u_2(g_2) = 1, \quad u_2(E \cup \{q'\}) \leq 0.8 + z \cdot 0.2 \cdot 1.$$

When $z = \frac{\sqrt{689}-17}{10} \approx 0.925$, the gap is $\min\left(\frac{u_1(g_2)}{u_1(E \cup \{q\})}, \frac{u_2(g_2)}{u_2(E \cup \{q'\})}\right) \geq \frac{5\sqrt{689}-115}{16} > 1.015$.

4

Approximate Efficiency under Chamberlin-Courant Scoring Rule

4.1 Introduction

Again we look at the multi-winner election problem in social choice. Recall that in this problem, the goal is to find a set of candidates (or winning committee) of fixed size k from voter preferences over the m candidates.

In this chapter, we consider the setting where the voters express *ordinal rankings* over these candidates, as we did in Chapter 2. In many applications such as parliamentary democracies, it is reasonable to assume voters can compare candidates or projects and hence can rank them ordinally, while they may not be able to articulate cardinal utilities for the same.

A classic method to aggregate the ordinal preferences is to use the (Borda)-Chamberlin-Courant scoring rule. It has been extensively studied in computational social choice [LB11, SFL16, SFS15, FLPT18, EFSS17, BFST19], starting with the work of Chamberlin and Courant [CC83].

In the minimization (or dissatisfaction) version, the Borda score of candidate c for voter v, denoted $r_v(c)$, is the ordinal rank of c in v's ranking. Here, the top-ranked candidate has Borda score 1, and the bottom-ranked candidate has score m.[1] Let $\mathcal{V}$ denote the set of all voters, with $n = |\mathcal{V}|$. Given a committee (that is, a set of candidates) T of size k, the Borda score of this committee is given by

$$r_{\mathcal{V}}(T) = \frac{1}{n} \sum_{v \in \mathcal{V}} \min_{c \in T} r_v(c). \tag{4.1}$$

Throughout the chapter, we will set the benchmark as the minimum possible score: $\textsc{Opt} := \min_{T \subseteq \mathcal{C}, |T| = k} r_{\mathcal{V}}(T)$.

The Chamberlin-Courant scoring rule assumes each voter is represented by the best candidate in T according to her ranking, so that optimizing this score achieves a form of aggregate efficiency. We will use the name Borda and Chamberlin-Courant interchangeably in the context of this chapter.

Our goal is to study the computational complexity (in m, n) of finding good committees according to the Borda score. In particular (though not exclusively), we focus on the analysis of greedy algorithms, which are appealing for their simplicity and ease of use, especially in settings involving human decision making, such as parliamentary elections or participatory budgeting with ordinal preferences.

4.1.1 Our Contribution

The Borda score is an ordinal version of the celebrated k-medians problem [AGK+04]. Unfortunately, for the ordinal version, it is not possible to approximate the minimum score, $\textsc{Opt}$, to any constant factor in polynomial time unless $\mathsf{P} = \mathsf{NP}$ [SFS15].

The GREEDY *Algorithm.* A natural algorithm for the 1-Borda score is the GREEDY algorithm that iteratively adds the candidate that decreases the 1-Borda score the

[1] Existing literature also uses a score of 0 for the best ranked and $m - 1$ for the worst ranked candidates. Since our results concern absolute scores, they carry over to this setting by simply subtracting 1 from the bounds. We use a minimum score of 1 since it is the more challenging setting for showing hardness results.

most. This algorithm was analyzed in [LB11] as follows. Consider the maximization (or satisfaction) version where the score of a candidate c for voter v is $m+1-r_v(c)$, so that the score of a committee T is $m+1-r_{\mathcal{V}}(T)$. Clearly, the maximization and minimization versions have the same optimum solutions, though they are very different from an approximation perspective. It is easy to check that the maximization objective is submodular [LB11], so that Greedy is a $\left(1-\frac{1}{e}\right)$-approximation by the classic result of [NWF78]. However, this analysis only shows that Greedy yields a solution of score at most m/e for the minimization objective.

Our first, and technically most challenging, contribution is an almost-tight analysis of this Greedy heuristic for the minimization version. In Section 4.3, we show that it achieves a score (given by Eq. (4.1)) of at most $2 \cdot \frac{m+1}{k+1}$ for any instance with m candidates from which we need to choose a committee of size k. We complement this analysis by exhibiting an instance where Greedy has score at least $1.962 \cdot \frac{m+1}{k+1}$.

For the maximization (or satisfaction) version, since the maximum possible score is m, the above result directly implies the following theorem.

Theorem 4.1.1. Greedy *is a* $\left(1-\frac{2}{k+1}\right)$*-approximation for the maximization version of* 1*-Borda score.*

For k larger than a small constant, this *significantly improves* the submodularity-based analysis [LB11] that only yields a $(1-1/e)$-approximation. Furthermore, it also improves on the best known approximation algorithm for this problem, Algorithm P in [SFS15], which achieves an approximation factor of $\left(1-O\left(\frac{\ln k}{k}\right)\right)$.

At a technical level, the standard analysis of Greedy for maximizing submodular functions shows that the next candidate yields an improvement in objective that is at least $1/k$ fraction of the gap between the current solution and the optimum. We use Cauchy-Schwarz inequality on per-voter improvements to show an overall improvement per step that has a *quadratic* dependence on the gap. This yields a significant improvement when the gap is large, and is the crux of why we are able to improve the upper-bound analysis of the maximization version significantly. Our lower bound instance works by carefully choosing per-voter improvements that make Cauchy-Schwarz inequality almost tight. This requires a non-trivial construction where the candidates chosen by GREEDY have ranks that lie on a carefully chosen spiral, and these are interspersed with candidates for whom voters' preferences are random. The ranks in each subsequent layer of the spiral decrease by a factor equal to the golden ratio.

A Benchmark and an Optimal Algorithm. The next natural question we ask is: How much better can we do in polynomial time? In Section 4.4, we show some hardness results for the minimization version (Eq (4.1)). Our main result *significantly improves* the constant factor hardness of approximation result of [SFS15] and shows the following:

Theorem 4.1.2. *Unless* $\mathsf{ZPP} = \mathsf{NP}$*, no polynomial-time algorithm can distinguish between instances with* $\textsc{Opt} \geq (1 - o(1)) \cdot \frac{m+1}{k+1}$ *from those with either:*

1. $\textsc{Opt} \leq \left(\frac{m+1}{k+1}\right)^{\delta}$*, where* $\delta \in (0, 1)$ *is a constant; or*

2. $\textsc{Opt} \leq \frac{1}{k^{\alpha}} \cdot \frac{m+1}{k+1}$*, where* $\alpha > 0$ *is a constant.*

This construction yielding this theorem is delicate. We require the full power of Feige's hardness proof of MAX COVER [Fei98], in particular, that it works on "regular" instances where each set has the same size, and where a collection of disjoint sets cover the instance completely in the "YES" case.

Theorem 4.1.2 motivates us to define the score $\frac{m+1}{k+1}$ as a reasonable benchmark for this problem, and we call any efficient algorithm achieving this score as an "optimal algorithm". Such a benchmark is appealing in that it helps us analyze other simple and natural algorithms that have been proposed in literature, and perform a more fine-grained comparison. As we have already seen, the GREEDY algorithm is always within a factor of 2 of this benchmark.

We now observe that if we pick a subset of k candidates at random from $\mathcal{C}$, the expected score is exactly the benchmark $\frac{m+1}{k+1}$. We therefore denote the score $\frac{m+1}{k+1}$ as RAND. Now, we can design a deterministic optimal algorithm via derandomizing this randomized algorithm. Interestingly, we show that this derandomization yields a greedy algorithm that is exactly the same as the BANZHAF algorithm proposed in [FLPT18] as a polynomial time heuristic for this problem. In that work, the BANZHAF algorithm was derived by viewing the problem as a cooperative game where players are candidates, and coalitions are committees, and adapting the notion of Banzhaf score of coalitions [Ban65, DS79]. It was emprically shown to be a very effective heuristic for this problem, beating GREEDY on most instances. We justify this empirical observation by viewing the BANZHAF algorithm instead as a derandomization of an optimal randomized algorithm.

In summary, we show the following theorem.

Theorem 4.1.3. *The* BANZHAF *algorithm achieves a minimization objective of at most* RAND $= \frac{m+1}{k+1}$ *in polynomial time, and is a* $\left(1 - \frac{1}{k+1}\right)$*-approximation to the maximization objective of Borda.*

To complete the picture, we show that an easy consequence of Theorem 4.1.2 is that the approximation factor of $\left(1 - \frac{1}{k+1}\right)$ is best possible for the maximization version unless $\mathsf{NP} = \mathsf{ZPP}$.

Committee Monotonicity. One appealing property of GREEDY is that it is *committee-monotone* [EFSS17]: The committee found for a smaller k is always a subset of a committee found for larger k's. This is immediate because GREEDY adds the next candidate to the committee based on the improvement in the 1-Borda score, and this improvement does not depend on k. On the other hand, the BANZHAF algorithm requires knowledge of k at each greedy step, and is therefore not committee-monotone. We therefore ask: Is there a committee-monotone algorithm that can achieve the benchmark RAND? In Section 4.5, we answer this question in the negative: There exist instances where any committee-monotone algorithm has score at least $1.015 \cdot$RAND. This shows a separation between committee-monotone algorithms and an optimal algorithm such as the BANZHAF algorithm.

4.2 Preliminaries

We consider the problem of selecting a subset of cardinality k from a set $\mathcal{C}$ of m candidates. We call this subset a *committee*. A set $\mathcal{V}$ of n voters express their preferences on the candidates ordinally. Each voter v has a bijective ranking function $r_v : \mathcal{C} \to \{1, 2, \ldots, m\}$, and v prefers those c's with smaller $r_v(c)$. For example, the top-ranked candidate of v, denoted by $c_{\mathrm{top}(v)}$, satisfies $r_v\left(c_{\mathrm{top}(v)}\right) = 1$, and the bottom-ranked $c_{\mathrm{bot}(v)}$ satisfies $r_v\left(c_{\mathrm{bot}(v)}\right) = m$.

In Borda score, the cost for a voter v of a committee T is her rank of the top candidate in T: $r_v(T) = \min_{c \in T} r_v(c)$. Further, the Borda score of a committee T is the average cost for all voters:

$$r_{\mathcal{V}}(T) = \frac{1}{n} \sum_{v \in \mathcal{V}} r_v(T) = \frac{1}{n} \sum_{v \in \mathcal{V}} \min_{c \in T} r_v(c).$$

Fix a voter v and look at her ranking on $\mathcal{C}$. If we pick a random size-k subset of $\mathcal{C}$, the t^{th} smallest rank is $t \cdot \frac{m+1}{k+1}$ in expectation. (See Section A.3 for a proof of this well-known fact.) Therefore, the expected performance of a random committee is

$$\mathop{\mathbf{E}}_{T \subseteq \mathcal{C}}[r_{\mathcal{V}}(T)] = \frac{1}{n} \sum_{v \in \mathcal{V}} \mathop{\mathbf{E}}_{T \subseteq \mathcal{C}}[r_v(T)] = \frac{m+1}{k+1}.$$

Define the benchmark Rand to be the expected performance of a random committee – Rand $:= \frac{m+1}{k+1}$. We will justify this benchmark in the subsequent sections.

We consider two simple committee-selection rules: Greedy and Banzhaf. These algorithms run in k iterations, during which they build sets $\varnothing = T_0 \subsetneq T_1 \subsetneq \cdots \subsetneq T_k$, and declare T_k as the selected committee.

In the j^{th} iteration, Greedy picks candidate $c_j \in \mathcal{C} \setminus T_{j-1}$ that minimizes $r_{\mathcal{V}}(T_{j-1} \cup \{c_j\})$, and let $T_j = T_{j-1} \cup \{c_j\}$. Banzhaf [Ban65, FLPT18], on the other hand, picks candidate $c_j \in \mathcal{C} \setminus T_{j-1}$ to minimize

$$\sum_{\substack{S \subseteq \mathcal{C}: |S|=k \\ S \supseteq T_{j-1} \cup \{c_j\}}} r_{\mathcal{V}}(S)$$

in the j^{th} iteration, and then sets $T_j = T_{j-1} \cup \{c_j\}$. In other words, it greedily picks the candidate that minimizes the final score if the rest of the committee is chosen uniformly at random. Both Greedy and Banzhaf can run in polynomial time [FLPT18].

Throughout the chapter, we use Rand, Greedy and Banzhaf to denote either the algorithms or their performances, which should be clear from the context.

4.3 Analysis of Greedy

In this section, we analyze the performance of Greedy, evaluated with respect to the benchmark Rand. We first show an upper bound that Greedy $\leq 2 \cdot$ Rand, and then present an almost-matching lower-bound instance where Greedy $> 1.962 \cdot$ Rand.

4.3.1 Upper Bound

Now we show Greedy $\leq 2 \cdot$ Rand as an upper bound. We first present the following lemma, which gives a lower bound on the improvement at each iteration.

Lemma 4.3.1. *Let T_t and T_{t+1} be the set of candidates produced by* Greedy *in the t^{th} and $(t+1)^{st}$ iterations, and $r_{\mathcal{V}}(T_t)$, $r_{\mathcal{V}}(T_{t+1})$ be their respective score. We have:*

$$r_{\mathcal{V}}(T_t) - r_{\mathcal{V}}(T_{t+1}) \geq \frac{\sum_{v\in\mathcal{V}} r_v(T_t)(r_v(T_t)-1)}{2n(m-t)}.$$

Proof. For a candidate $c \notin T_t$, define $\Delta_c := r_{\mathcal{V}}(T_t) - r_{\mathcal{V}}(T_t \cup \{c\})$, i.e., the current marginal contribution of c to the Borda score. Taking the sum of Δ_c over $c \notin T_t$:

$$\sum_{c\in\mathcal{C}\setminus T_t} \Delta_c = \frac{1}{n}\sum_{v\in\mathcal{V}}\sum_{j=1}^{r_v(T_t)-1} j = \frac{\sum_{v\in\mathcal{V}} r_v(T_t)(r_v(T_t)-1)}{2n}.$$

Greedy chooses $c^* = \arg\max_c \Delta_c$ at the $(t+1)^{\text{st}}$ iteration, giving us

$$r_{\mathcal{V}}(T_t) - r_{\mathcal{V}}(T_{t+1}) = \Delta_{c^*} \geq \frac{1}{m-t}\sum_{c\in\mathcal{C}\setminus T_t} \Delta_c = \frac{\sum_{v\in\mathcal{V}} r_v(T_t)(r_v(T_t)-1)}{2n(m-t)}. \qquad \square$$

Now we prove our upper bound of 2.

Theorem 4.3.2. Greedy $\leq 2 \cdot$ Rand.

Proof. We prove by induction. As the base case where $k = 1$, Greedy $\leq m < m+1 = 2 \cdot$ Rand. Now suppose that the claim holds for some $k-1$ and we will prove that it also holds for k. By induction hypothesis, we have:

$$r_{\mathcal{V}}(T_{k-1}) \leq 2 \cdot \frac{m+1}{k}.$$

If $r_{\mathcal{V}}(T_{k-1}) \leq 2 \cdot \frac{m+1}{k+1}$, then $r_{\mathcal{V}}(T_k) \leq r_{\mathcal{V}}(T_{k-1}) \leq 2 \cdot \frac{m+1}{k+1}$ finishes the proof. Thus, we only need to consider the following case:

$$2 \cdot \frac{m+1}{k+1} < r_{\mathcal{V}}(T_{k-1}) \leq 2 \cdot \frac{m+1}{k}.$$

We now have the following, where the first inequality is by Lemma 4.3.1 and second by Cauchy-Schwarz inequality:

$$\begin{aligned} r_{\mathcal{V}}(T_{k-1}) - r_{\mathcal{V}}(T_k) &\geq \frac{\sum_{v \in \mathcal{V}} r_v(T_{k-1})(r_v(T_{k-1}) - 1)}{2n(m-k+1)} \\ &\geq \frac{\frac{1}{n}(\sum_{v \in \mathcal{V}} r_v(T_{k-1}))^2 - \sum_{v \in \mathcal{V}} r_v(T_{k-1})}{2n(m-k+1)} \\ &= \frac{(\sum_{v \in \mathcal{V}} r_v(T_{k-1}))^2}{2n^2(m+1)} \cdot \frac{m+1}{m-k+1} \cdot \frac{\sum_{v \in \mathcal{V}}(r_v(T_{k-1}) - 1)}{\sum_{v \in \mathcal{V}} r_v(T_{k-1})}. \end{aligned}$$

Since $r_{\mathcal{V}}(T_{k-1}) \geq 2 \cdot \frac{m+1}{k+1}$ by assumption, we have:

$$\begin{aligned} \frac{m+1}{m-k+1} \cdot \frac{\sum_{v \in \mathcal{V}}(r_v(T_{k-1}) - 1)}{\sum_{v \in \mathcal{V}} r_v(T_{k-1})} &\geq \frac{m+1}{m-k+1} \cdot \frac{2 \cdot \frac{m+1}{k+1} - 1}{2 \cdot \frac{m+1}{k+1}} \\ &= \frac{2(m+1) - k - 1}{2(m+1) - 2k} \geq 1. \end{aligned}$$

Combining the previous two inequalities, we therefore have:

$$r_{\mathcal{V}}(T_{k-1}) - r_{\mathcal{V}}(T_k) \geq \frac{(\sum_{v \in \mathcal{V}} r_v(T_{k-1}))^2}{2n^2(m+1)} = \frac{r_{\mathcal{V}}^2(T_{k-1})}{2(m+1)},$$

which is equivalent to:

$$r_{\mathcal{V}}(T_k) \leq -\frac{1}{2(m+1)} r_{\mathcal{V}}^2(T_{k-1}) + r_{\mathcal{V}}(T_{k-1}).$$

Notice that the right hand side is a quadratic function in $r_{\mathcal{V}}(T_{k-1})$, which is monotonically increasing for $r_{\mathcal{V}}(T_{k-1}) \leq m+1$. Since $r_{\mathcal{V}}(T_{k-1}) \leq 2 \cdot \frac{m+1}{k} \leq m+1$, the right hand side reaches its maximum at $2 \cdot \frac{m+1}{k}$. Thus, we have:

$$r_{\mathcal{V}}(T_k) \leq -\frac{1}{2(m+1)} \cdot \left(\frac{2(m+1)}{k}\right)^2 + \frac{2(m+1)}{k} \leq \frac{2(m+1)}{k+1},$$

which concludes our induction. $\square$

Proof of Theorem 4.1.1. For the maximization version, the above result implies GREEDY achieves score at least $(m+1) \cdot \left(1 - \frac{2}{k+1}\right)$. Since the maximum possible score is m, this implies that GREEDY is a $\left(1 - \frac{2}{k+1}\right)$-approximation.

4.3.2 Lower Bound

Now we complement our result with a lower-bound example for GREEDY.

Theorem 4.3.3. *There exists an instance in which* $r_{\mathcal{V}}(T_k) > 1.962 \cdot \text{RAND}$.

Construction. In the sequel, we will prove the above theorem. In the instance we construct, m, n, and k are all sufficiently large. For convenience of illustration, we scale down the ranks by a factor of m: now the ranks are $\frac{1}{m}, \frac{2}{m}, \ldots, \frac{m-1}{m}, 1$. As $m \to \infty$, $\frac{1}{m} \to 0$, so the set of ranking $\{\frac{1}{m}, \frac{2}{m}, \ldots, 1\}$ will become dense in $[0, 1]$, and thus we regard the ranking as being continuous from 0 to 1. Our goal becomes to construct an instance in which GREEDY gives $r_{\mathcal{V}}(T_k) > 1.962 \cdot \frac{1}{k+1}$.

There are sufficiently many voters, enabling us to view them as a continuum from 0 to 1, forming a circle (the base in Fig. 4.1) with angular position ranging from 0 to 2π. Imagine that each voter writes down her favorite, her second favorite, . . . , her least favorite candidate in that order vertically. The result is the side of a cylinder with height 1, as depicted in Fig. 4.1. Each point on the side identifies a candidate, whose distance to the top, d, indicates the corresponding voter ranks him as her $(dm)^{\text{th}}$ favorite candidate (i.e., the candidate has a rank of d in the voter's preference after scaling).

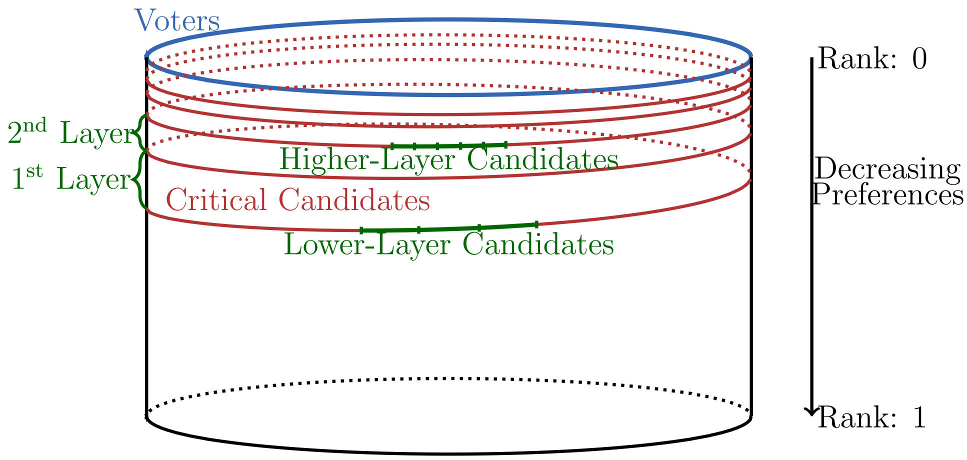

FIGURE 4.1: Construction of the Bad Instance for GREEDY

We divide the set of candidates into two types – *critical* and *dummy*. The former set has size $k \ll m$, and the latter has size $m-k$. Our proof will show that GREEDY will choose the critical candidates in a fixed order, and will not choose any dummy candidate.

The critical candidates are present in ℓ "layers" as shown in the red spiral in Fig. 4.1, where ℓ is sufficiently large. This figure shows the ranks of the critical candidates in the voters' profiles. We parametrize this spiral by θ, which maps to the voter at the corresponding angular position $2\pi\theta$. We place critical candidates in order, where each candidate appears a number of times consecutively on the spiral. Therefore, each voter has one critical candidate from each layer $t = 0, 1, \ldots, \ell$ in the spiral part of its ranking.

In the t^{th} layer, the parameter θ lies in $[t-1, t)$. The critical candidate when the parameter is θ has rank $g(\theta) = a\varphi^{\theta}$ for the voter at angular position $2\pi\theta$. Here, φ denotes the golden ratio $\frac{\sqrt{5}-1}{2} \approx 0.618$, and a is a sufficiently small constant so that rounding to the nearest integer does not change the analysis. This critical candidate is placed for a certain length $h(\theta)$ on the spiral, which means this candidate appears at rank $g(\theta)$ for voters in the range $[2\pi\theta, 2\pi(\theta + h(\theta))]$. In our construction, $h(\theta)$ will be very small, so that we will say this candidate appears $h(\theta)$ times at rank $g(\theta)$ for parameter θ. The greater θ is, the smaller $h(\theta)$ has to be, and we will calculate its expression later.

For the convenience of analysis, at the layer $t = 0$, that is, for $\theta \in [-1, 0)$, there is a special candidate appearing on the spiral throughout the layer. This special candidate is picked first by GREEDY. Other than its appearance on the spiral, any critical candidate is placed at the very bottom, i.e., rank 1, for the other voters. Denote the total number of critical candidates by k. Then we have $m - k$ dummy candidates. These dummy candidates are symmetrically placed at other ranks. We copy each voter $(m - k)!$ times, once for each possible permutation of the dummy candidates to place in the remaining ranks.

The idea of this construction is to trick GREEDY into picking every critical candidate on the spiral in order, while in fact, lower-layer critical candidates have no contribution to the objective once higher-layer ones have been selected. The following analysis computes the optimal parameters to realize this plan.

Not Choosing a Dummy Candidate. We first ensure GREEDY does not choose a dummy candidate in this instance by setting $h(\theta)$ properly. We assume that GREEDY chooses critical candidates in increasing order of θ, and we will justify this assumption later.

To simplify notation, denote $X = \int_0^1 a\varphi^\theta \, \mathrm{d}\theta$ and $Y = \int_0^1 a^2\varphi^{2\theta} \, \mathrm{d}\theta$. Computing these explicitly:

$$X = \frac{a}{\ln \varphi}(\varphi - 1), \qquad Y = \frac{a^2}{2\ln\varphi}(\varphi^2 - 1) = X^2 \frac{(\varphi + 1)\ln\varphi}{2(\varphi - 1)}.$$

Using this notation, consider the critical candidate at the beginning of the first layer, that is, at $\theta = 0$. Since GREEDY chooses the candidate at layer $t = 0$, the decrease in score due to this critical candidate is:

$$h(0) \cdot (g(-1) - g(0)) = h(0) \cdot a \cdot \left(\frac{1}{\varphi} - 1\right) = h(0) \cdot a \cdot \varphi. \tag{4.2}$$

where we have used that since φ is the golden ratio, $\varphi + \varphi^2 = 1$.

Now consider the dummy candidates. Just after GREEDY has chosen the special candidate at layer $t = 0$, each such candidate improves the rank of $g(\theta - 1)$ fraction of voters at $\theta \in [0, 1)$. This is because we placed all permutations of the dummy candidates with each voter θ, and GREEDY has already chosen the special candidate. By the same reasoning, conditioned on improvement, the average improvement is $g(\theta - 1)/2$. Therefore, the decrease in score due to a dummy candidate is:

$$\int_0^1 \frac{g^2(\theta - 1)}{2} \, \mathrm{d}\theta = \frac{a^2}{2\varphi^2} \int_0^1 \varphi^{2\theta} \, \mathrm{d}\theta = \frac{1}{2\varphi^2} \cdot Y. \tag{4.3}$$

Since we want GREEDY to choose the critical candidate, we need to set

$$h(0) = \frac{Y}{2\varphi^3 a}.$$

By the symmetry of the spiral, an identical calculation now holds for all $\theta > 0$. To make GREEDY choose the critical candidate at this location (assuming it has chosen critical candidates for smaller values of θ), we need:

$$h(\theta) = \frac{Y}{2\varphi^3 a} \varphi^\theta.$$

Note that $h(\theta)$ depends linearly on a, so that for very small a, we can pretend this set of voters lies exactly at θ. Further, $h(\theta)$ is decreasing with θ.

Choosing Critical Candidates in Order. We now show that GREEDY chooses the critical candidates following the order on the spiral.

Lemma 4.3.4. Greedy *chooses the critical candidates in increasing order of θ.*

Proof. The calculation is identical at any step of Greedy, so we focus on the step where Greedy is at the beginning of the first layer, that is, considering the critical candidate at $\theta = 0$. Recall that Greedy has chosen the special candidate at layer $t = 0$. The previous analysis showed that the critical candidate at $\theta = 0$ yields decrease of $\frac{Y}{2\varphi^2}$. For critical candidates in the same layer $t = 1$ (that is, for $\theta \in [0, 1)$), the contribution of the candidate at θ is

$$h(\theta) \cdot (g(\theta - 1) - g(\theta)) = \frac{Y}{2\varphi^3}\varphi^\theta \left(\varphi^{\theta-1} - \varphi^\theta\right) = \frac{Y}{2\varphi^2} \cdot \varphi^{2\theta},$$

which decreases with θ, so that the current candidate, $\theta = 0$, offers the best decrease. Here, we have used that since φ is the golden ratio, $\varphi^2 + \varphi = 1$.

For $t \geq 1$, suppose we instead considered a candidate $t + \theta$ for $\theta \in [0, 1)$ located in layer $t + 1$. Conditioned on having chosen layer $t = 0$, this candidate gives a contribution of

$$\begin{aligned} h(t+\theta) \cdot (g(\theta - 1) - g(t+\theta)) &\leq h(t) \cdot (g(-1) - g(t)) \\ &\leq \max\big(h(2) \cdot g(-1),\ h(1) \cdot (g(-1) - g(1))\big) \\ &= \max\left(\frac{Y}{2\varphi^2},\ \frac{Y}{2\varphi^2 a} \cdot a\left(\frac{1}{\varphi} - \varphi\right)\right) = \frac{Y}{2\varphi^2}, \end{aligned}$$

where the first inequality uses that $h(\theta)$ is decreasing in θ, and that $\varphi < 1$.

Therefore, Greedy will pick the critical candidate at $\theta = 0$ instead of another candidate at the same or a higher layer. Since the argument is identical at each θ, Greedy picks critical candidates in order on the spiral. □

The Lower Bound. So far we have shown that Greedy chooses critical candidates in increasing order of layers and does not choose dummy candidates. We finally put it all together and show the following bound, which completes the proof of Theorem 4.3.3.

Proof of Theorem 4.3.3. The number of critical candidates on the t^{th} layer ($\theta \in [t-1, t)$) is

$$\int_{t-1}^{t} \frac{1}{h(\theta)} \,\mathrm{d}\theta = \frac{2\varphi^3 a}{Y} \int_{t-1}^{t} \varphi^{-\theta} \,\mathrm{d}\theta = \frac{2a}{\varphi^{t-3} Y} \int_{-1}^{0} \varphi^{-\theta} \,\mathrm{d}\theta = \frac{2a}{\varphi^{t-3} Y} \int_{0}^{1} \varphi^{\theta} \,\mathrm{d}\theta = \frac{2X}{\varphi^{t-3} Y}.$$

Therefore, when it is done with the ℓ^{th} layer, the number of candidates Greedy has picked is

$$k = \frac{2X}{\varphi^{\ell-3} Y}(1 + \varphi + \varphi^2 + \cdots + \varphi^{\ell-1}) \to \frac{2X}{(1-\varphi)\varphi^{\ell-3} Y}$$

when ℓ is large. Meanwhile, the Borda score of Greedy is

$$r_{\mathcal{V}}(T_k) = \int_0^1 g(\ell - 1 + \theta) \,\mathrm{d}\theta = \varphi^{\ell-1} X.$$

Therefore, the approximation ratio is

$$\begin{aligned} (k+1) r_{\mathcal{V}}(T_k) &\geq \frac{2X}{(1-\varphi)\varphi^{\ell-3} Y} \cdot \varphi^{\ell-1} X \\ &= \frac{2\varphi^2 X^2}{(1-\varphi) Y} \\ &= \frac{2\varphi^2}{(1-\varphi)} \cdot \frac{2(\varphi - 1)}{(\varphi+1)\ln\varphi} \\ &= -\frac{4\varphi^2}{(\varphi+1)\ln\varphi} > 1.962. \end{aligned}$$

□

4.4 Hardness and an Optimal Deterministic Algorithm

We justify our choice of benchmark Rand $= \frac{m+1}{k+1}$, and show that a deterministic algorithm, Banzhaf, achieves this benchmark optimally. First, notice that if the input consists of one voter for each possible preference of m candidates (thus $n = m!$), picking any committee has the same Borda score by symmetry, so Opt is just Rand. Thus, we have the following proposition:

Proposition 4.4.1. *For any m and k, there exist instances where* Opt = Rand.

4.4.1 Hardness Result: Theorem 4.1.2

We now show Theorem 4.1.2: Even if OPT is very small, it is computationally hard to significantly beat RAND. To prove this hardness result, we show a reduction from the decision version of the REGULAR MAX K-COVER problem.

Definition 4.4.2. In REGULAR MAX K-COVER, there is a universe U of n elements $\{a_1, a_2, \ldots, a_n\}$, and a family $\mathcal{F} = \{S_1, S_2, \ldots, S_z\}$ of subsets of U. Each S_i has the same size $\frac{n}{k}$. The *value* of an instance is the maximum size of the union of k sets from $\mathcal{F}$. For any constant $\varepsilon > 0$, we consider the following decision version:

- "YES" instances are those with value n. Therefore, there exist k disjoint sets each of size n/k that cover all the elements.
- "NO" instances are those with value at most $\frac{3}{4}n$.

The above problem known to be NP-hard to approximate via the following lemma that is implicit in the proof of Theorem 5.3 in [Fei98].

Lemma 4.4.3 ([Fei98])**.** *The decision version of* REGULAR MAX K-COVER *from Definition 4.4.2 is* NP-hard*, that is, unless* P = NP*, there is no polynomial time algorithm that can decide always answers "YES" for "YES" instances and answers "NO" for "NO" instances.*

Note that if the instance has value n, there exist k disjoint sets each of size n/k that cover all the elements. This aspect will be crucial in our reduction. Also needed in our reduction, we state the following lemma for constructing a profile with polynomial number of voters, where the best solution with score OPT has similar performance as RAND.

Lemma 4.4.4. *Fix any $\varepsilon > 0$ and let $n \geq \left\lceil \frac{m(k+1)^2}{\varepsilon^2} \right\rceil$. Consider the instance where the preference of each voter is an independent and uniformly random permutation.*

Let $\textsc{Opt}'$ *denote the expected value of the optimum score, and* $\textsc{Rand}' = \frac{m+1}{k+1}$*, then* $\Pr[\textsc{Opt} \leq (1-\varepsilon)\cdot \textsc{Rand}] < \frac{1}{2}$*, where the probability is over the randomness in the permutations.*

Proof. Fix any committee T of size k. Notice that $\mathbf{E}[r_{\mathcal{V}}(T)] = \textsc{Rand}'$ since the preferences are uniformly random. We have

$$\Pr[r_{\mathcal{V}}(T) - \textsc{Rand}' \leq -\varepsilon \cdot \textsc{Rand}'] = \Pr\left[\frac{1}{n}\sum_{v\in\mathcal{V}} r_v(T) - \textsc{Rand}' > \varepsilon\cdot\textsc{Rand}'\right]$$
$$\leq \exp\left(\frac{-2n(\varepsilon\cdot\textsc{Rand}')^2}{m^2}\right) \leq \exp\left(\frac{-2n\varepsilon^2}{(k+1)^2}\right),$$

where the second step comes from Hoeffding's inequality. By union bound,

$$\Pr[\textsc{Opt}' \leq (1-\varepsilon)\cdot\textsc{Rand}'] \leq \binom{m}{k}\cdot \Pr[r_{\mathcal{V}}(T) - \textsc{Rand}' \leq -\varepsilon\cdot\textsc{Rand}']$$
$$\leq \exp\left(\frac{-2n\varepsilon^2}{(k+1)^2} + m\right) \leq e^{-m} < \frac{1}{2}. \qquad \square$$

Now we are ready to prove Theorem 4.1.2.

Proof of Theorem 4.1.2. Fix a $\varepsilon > 0$ and let $\varepsilon' = 10\varepsilon$. We will choose ε appropriately later. Given any instance of Regular Max K-Cover with n elements and z sets each of size n/k (as in Definition 4.4.2), we construct the following instance for our problem:

- There are $N = nR$ voters v_{ij} where $i \in [n]$ and $j \in [R]$. We have $m = \frac{2}{\varepsilon'}kz$ candidates. The first z candidates $\{c_1, c_2, \ldots, c_z\}$ are "critical" candidates, and the other $m - z$ candidates are "dummy" candidates. Each voter corresponds to an element in the universe and each critical candidate corresponds to a set in Regular Max K-Cover.

- If a set S_i covers a_j, then voters v_{ij} for $j \in [R]$ rank c_j within top ε' fraction. Otherwise, v_{ij}'s rank c_j within bottom ε' fraction.

- Independently for each voter, fill the rest of her preferences with the $m - z$ dummy candidates uniformly randomly.

- The copies of a voter only differ in the ranking of the dummy candidates. We set the number of copies to be $R = \left\lceil \frac{10mk^2}{n\varepsilon^2} \right\rceil$. These copies are there to ensure Lemma 4.4.4 applies to the dummy candidates.

Clearly, the above construction has size $\text{poly}(1/\varepsilon, n, z, k)$. Let $\textsc{Opt}$ denote the optimal score on this instance. Recall that $\textsc{Rand} = \frac{m+1}{k+1}$. First suppose the instance of Regular Max K-Cover has value n ("YES" instance) so that there are k sets that cover all n elements, then it is easy to check that choosing the corresponding critical candidates as the committee yields $\textsc{Opt} \leq z < \varepsilon' \cdot \textsc{Rand}$.

On the other hand, suppose the instance of Regular Max K-Cover is such that any collection of sets of size k only covers at most $(1 - 1/e + \varepsilon)n \leq \frac{3}{4}n$ elements ("NO" instance). Consider any committee T and suppose $T = R \cup D$ where R is a subset of critical candidates and D is a subset of dummy candidates. Let $r = |R|$ and $d = |D| = k - r$. Let n' be the number of elements R covers in the Regular Max K-Cover instance. By assumption, $n - n' \geq \frac{n}{4}$ since any collection R covers at most $\frac{3}{4}n$ elements. Further, since the instance is regular, $n - n' \leq \frac{n}{k}r$, so that $n - n' \geq d\frac{n}{k}$.

Using Lemma 4.4.4, with probability $> \frac{1}{2}$ over the choice of the ranking of the dummy candidates, the optimal score of D on the $(n - n')R$ uncovered voters using the $m - z$ dummy candidates is greater than $(1 - \varepsilon)\frac{m-z}{d+1}$. Inserting the critical candidates cannot decrease this score for these voters, since the candidates in R appear last in their ordering. Further, we have assumed $m = \frac{2}{\varepsilon'}kz$. Therefore, with probability $> \frac{1}{2}$, we have:

$$\textsc{Opt} > \frac{n - n'}{n}(1 - \varepsilon)\frac{m - z}{d + 1} \geq \frac{n - n'}{n} \cdot \left(1 - \frac{\varepsilon'}{4}\right) \cdot \frac{m + 1}{d + 1}$$

We now split the analysis into two cases:

1. Suppose $d+1 \le \frac{k+1}{4}$. Since $n - n' \ge \frac{n}{4}$, we have

$$\textsc{Opt} > \frac{1}{4} \cdot \left(1 - \frac{\varepsilon'}{4}\right) \cdot \frac{m+1}{d+1} \ge \left(1 - \frac{\varepsilon'}{4}\right) \cdot \frac{m+1}{k+1} \ge (1-\varepsilon')\textsc{Rand}.$$

2. Suppose $d+1 \ge \frac{k+1}{4}$. Since $n - n' \ge \frac{d}{k}n$, and since $d, k = \omega(1)$, we have:

$$\textsc{Opt} > \frac{d}{k} \cdot \left(1 - \frac{\varepsilon'}{4}\right) \cdot \frac{m+1}{d+1} = \frac{d}{d+1}\frac{k+1}{k}\left(1 - \frac{\varepsilon'}{4}\right) \cdot \frac{m+1}{k+1} \ge (1-\varepsilon')\textsc{Rand}.$$

Therefore, our construction ensures that with probability $> \frac{1}{2}$, we have $\textsc{Opt} \ge (1-\varepsilon')\textsc{Rand}$ if the original Regular Max K-Cover instance has value at most $\frac{3}{4}n$.

Now suppose there is a polynomial time algorithm that can distinguish between instances with $\textsc{Opt} \le \varepsilon'\textsc{Rand}$ and $\textsc{Opt} \ge (1-\varepsilon')\textsc{Rand}$. Then, feeding the output of the above construction to this algorithm implies a coRP algorithm for the decision version of Regular Max K-Cover, which by Theorem 4.4.3 implies $\mathsf{NP} \subseteq \mathsf{coRP}$. Since $\mathsf{RP} \subseteq \mathsf{NP}$, this implies $\mathsf{RP} \subseteq \mathsf{coRP}$, so that $\mathsf{ZPP} = \mathsf{RP} \cap \mathsf{coRP} = \mathsf{RP}$. Since ZPP is symmetric with respect to "YES" and "NO" instances, this implies $\mathsf{ZPP} = \mathsf{coRP}$, so that $\mathsf{ZPP} = \mathsf{NP}$.

We now show how to set ε. For the first part of the theorem, we set $\varepsilon = \frac{1}{10}\left(\frac{k}{m}\right)^{1-\delta}$. This can be achieved by choosing m such that $\left(\frac{m}{k}\right)^{\delta} = 20kz$. Note that this ensures $m = \mathrm{poly}(k, z)$ when δ is a constant, so that the construction runs in polynomial time. For this setting, we have $\varepsilon' \cdot \frac{m+1}{k+1} \le \left(\frac{m+1}{k+1}\right)^{\delta}$, while $\varepsilon' = \left(\frac{k}{m}\right)^{1-\delta} = \left(\frac{1}{20kz}\right)^{\frac{1-\delta}{\delta}} = o(1)$, completing the proof.

For the second part of the theorem, we set $m = 20k^{1+\alpha}z$, and $\varepsilon = \frac{1}{10}\frac{1}{k^{\alpha}}$. Again, we have $m = \mathrm{poly}(k, z)$, and $\varepsilon' = o(1)$, completing the proof. $\square$

Theorem 4.1.2 now implies the following easy corollaries.

Corollary 4.4.5. *Unless* $\mathsf{NP} = \mathsf{ZPP}$*, there is no* k^α*-approximation to the Borda score for any constant* $\alpha > 0$*. Similarly, there is no* $\left(\frac{m+1}{k+1}\right)^{1-\delta}$*-approximation for any constant* $\delta \in (0, 1)$*.*

The next corollary adapts the hardness proof to the maximization version of the problem.

Corollary 4.4.6. *For the maximization version of Borda, there is no polynomial time* $\left(1 - \frac{1-\varepsilon}{k+1}\right)$*-approximation for constant* $\varepsilon \in (0, 1/2)$ *unless* $\mathsf{NP} = \mathsf{ZPP}$*.*

Proof. Set $\varepsilon > 0$ to be a small constant in the proof of Theorem 4.1.2. Then, in the "NO" instance, the maximization score is at most $(m + 1)\left(1 - \frac{1-\varepsilon}{k+1}\right)$, while for the "YES" instance, the score is at least $(m + 1)\left(1 - \frac{\varepsilon}{k+1}\right)$. For $\varepsilon \in (0, 1/2)$, the approximation factor achievable is therefore at most $\left(1 - \frac{1-2\varepsilon}{k+1}\right)$, completing the proof. $\square$

4.4.2 An Optimal Deterministic Algorithm

Given the lower bound and the hardness result, an immediate question is whether there is a deterministic rule to achieve the benchmark Rand. We answer in the affirmative: The Banzhaf algorithm [Ban65, DS79, FLPT18] can be viewed as a derandomization of Rand: Instead of randomly picking a candidate at each iteration, it picks the candidate that gives the best expected performance if the rest of the committee is randomly constructed. It is shown in [FLPT18] that this algorithm runs in polynomial time. The following theorem implies Theorem 4.1.3.

Theorem 4.4.7. Banzhaf $\leq$ Rand*.*

Proof. Recall that Banzhaf builds sets $\varnothing = T_0 \subsetneq T_1 \subsetneq \cdots \subsetneq T_k$, where at step j,

$\textsc{Banzhaf}$ picks $c_j \in \mathcal{C} \setminus T_{j-1}$ such that:

$$c_j = \text{argmin}_{c \in \mathcal{C} \setminus T_{j-1}} \sum_{\substack{S \subseteq \mathcal{C} : |S| = k \\ S \supseteq T_{j-1} \cup \{c\}}} r_{\mathcal{V}}(S). \tag{4.4}$$

We now use induction to show that for any $j \in \{0, 1, \ldots, k\}$,

$$\frac{1}{\binom{m-j}{k-j}} \sum_{\substack{S \subseteq \mathcal{C} : |S| = k \\ S \supseteq T_j}} r_{\mathcal{V}}(S) \leq \textsc{Rand},$$

which is clearly true when $j = 0$, and gives the desired result $\textsc{Banzhaf} \leq \textsc{Rand}$ when $j = k$.

For the inductive step, assume it holds for some $j-1$. Now in the j^{th} iteration, we have the following inequalities that complete the proof. Here, the first step follows since $\textsc{Banzhaf}$ picks c_j in step j. The second step follows since $\textsc{Banzhaf}$ solves Eq (4.4), so that the score from adding c_j beats the average score of adding one of the $m - j + 1$ candidates in $C \setminus T_{j-1}$. The final equality follows since $\binom{m-j+1}{k-j+1} = \frac{m-j+1}{k-j+1}\binom{m-j}{k-j}$, and by observing that for any $S \supseteq T_{j-1}$, there are $k - j + 1$ choices of $c \in S \setminus T_{j-1}$.

$$\begin{aligned}
\frac{1}{\binom{m-j}{k-j}} \sum_{\substack{S \subseteq \mathcal{C} : |S| = k \\ S \supseteq T_j}} r_{\mathcal{V}}(S) &= \frac{1}{\binom{m-j}{k-j}} \sum_{\substack{S \subseteq \mathcal{C} : |S| = k \\ S \supseteq T_{j-1} \cup \{c_j\}}} r_{\mathcal{V}}(S) \\
&\leq \frac{1}{\binom{m-j}{k-j}} \cdot \left(\frac{1}{m-j+1} \sum_{c \in C \setminus T_{j-1}} \sum_{\substack{S \subseteq \mathcal{C} : |S| = k \\ S \supseteq T_{j-1} \cup \{c\}}} r_{\mathcal{V}}(S) \right) \\
&= \frac{1}{\binom{m-j}{k-j}} \cdot \frac{k-j+1}{m-j+1} \cdot \left(\frac{1}{k-j+1} \sum_{c \in C \setminus T_{j-1}} \sum_{\substack{S \subseteq \mathcal{C} : |S| = k \\ S \supseteq T_{j-1} \cup \{c\}}} r_{\mathcal{V}}(S) \right) \\
&= \frac{1}{\binom{m-j+1}{k-j+1}} \sum_{\substack{S \subseteq \mathcal{C} : |S| = k \\ S \supseteq T_{j-1}}} r_{\mathcal{V}}(S) \leq \textsc{Rand}.
\end{aligned}$$

$\square$

To complete the proof of Theorem 4.1.3, for the maximization objective, Banzhaf achieves a value at least $(m+1)\left(1-\frac{1}{k+1}\right)$. Since the maximum possible value is m, this implies a $\left(1-\frac{1}{k+1}\right)$-approximation.

4.5 Lower Bound on Committee-Monotone Algorithms

A nice property of Greedy is that it is committee-monotone: In each iteration, the candidate chosen by Greedy only depends on which candidates have been chosen in previous iterations and not on k, and thus when k increases, the committee selected by Greedy includes all the candidates Greedy used to select. On the other hand, Banzhaf does not satisfy committee monotonicity, as the candidates chosen by Banzhaf does depend on k.

This naturally brings up the question: Is there a committee-monotone algorithm which is optimal with respect to the benchmark Rand? We answer this question in the negative, by presenting a lower bound of 1.015 for all committee-monotone algorithms.

Theorem 4.5.1. *For any large enough m, there exist instances with m candidates where any committee-monotone algorithm* Alg *satisfies $r_{\mathcal{V}}(T_k) > 1.015 \cdot$* Rand *for some value $k \in \{1, 2\}$. Here, T_k is the set of candidates* Alg *chooses when the size of this set is k.*

Proof. The construction goes as follows: There are two types of candidates, X and Y. Candidates of type X are ranked between $[am, bm]$ by every voter and candidates of type Y are ranked between $[1, am] \cup [bm, m]$ by every voter, where $0 < a < b < 1$ are two parameters. We construct sufficiently many voters so that all candidates of the same type are symmetric (by having all permutations of candidates of type X and those of type Y). We want to find proper a and b, so that when $k = 1$, the optimal choice is to choose a candidate of type X, while when $k = 2$, the optimal choice is to choose two candidates both of type Y. This means that no committee-monotone algorithm can produce optimal choice for both the first iteration and the second iteration. We optimize over a and b to find the maximum lower bound.

In particular, the search procedure goes as follows. Let $r_{\mathcal{V}}(Y)$ denote the 1-Borda score of choosing a candidate of type Y; $r_{\mathcal{V}}(XX)$ denote the score of choosing two candidates both of type X; and $r_{\mathcal{V}}(XY)$ denote the score of choosing a candidate of type X and a candidate of type Y. We can easily see that, when m goes to infinity, up to an $o(1)$ additive error:

$$\begin{cases} \frac{1}{m+1} r_{\mathcal{V}}(Y) = \frac{a}{2} \cdot \Pr[Y \text{ is at } [1, am]] + \frac{1+b}{2} \cdot \Pr[Y \text{ is at } [bm, m]] = \frac{a}{2} \cdot \frac{a}{1-(b-a)} + \frac{1+b}{2} \cdot \frac{1-b}{1-(b-a)} \\ \frac{1}{m+1} r_{\mathcal{V}}(XX) = \frac{2a+b}{3} \\ \frac{1}{m+1} r_{\mathcal{V}}(XY) = \frac{a}{2} \cdot \Pr[Y \text{ is at } [1, am]] + \frac{a+b}{2} \cdot \Pr[Y \text{ is at } [bm, m]] = \frac{a}{2} \cdot \frac{a}{1-(b-a)} + \frac{a+b}{2} \cdot \frac{1-b}{1-(b-a)} \end{cases}.$$

A committee-monotone algorithm either chooses Y in the first iteration, or chooses XX or XY in the first two iterations. Therefore, to give a lower bound, we maximize $\min\left(\frac{2}{m+1} r_{\mathcal{V}}(Y), \frac{3}{m+1} r_{\mathcal{V}}(XX), \frac{3}{m+1} r_{\mathcal{V}}(XY)\right)$ (note that the value on the numerator corresponds to the value of $k+1$) over $0 < a < b < 1$, and find that, for $a = 0.377$ and $b = 0.552$, it achieves a lower bound greater than 1.015. $\square$

5

Approximate Efficiency in Bilateral Trade

5.1 Introduction

In a bilateral trade, a seller holds an item and is trying to sell it to a buyer. The buyer's private value v is drawn from a cumulative distribution function (CDF) F, and the seller's private cost c for selling the item is independently drawn from a CDF G. If a transaction happens between v and c, the society as a whole gains utility of $v - c$. The *gains-from-trade* (GFT) refer to the expected utility gain from trading. If trading probability between v and c is $x(v, c)$, then

$$\mathsf{GFT} = \mathop{\mathbf{E}}_{\substack{v \sim F \\ c \sim G}} \big[(v - c) \cdot x(v, c)\big].$$

Ideally, to maximize GFT, a trade should always happen when $v > c$ and never happen when $v < c$. The resulting optimal GFT is termed the *first best* (FB). Namely,

$$\mathsf{FB} = \mathop{\mathbf{E}}_{\substack{v \sim F \\ c \sim G}} \big[(v - c) \cdot \mathbb{1}\{v \geq c\}\big].$$

The seminal work of [MS83] shows that if both agents are self-interested, it is impossible to devise a Bayesian incentive-compatible (BIC), individually rational

(IR) and weakly budget balanced[1] (WBB) mechanism that achieves the first-best GFT, as long as the distribution supports of F and G "overlap". This impossibility result motivates the natural question of whether there exists a mechanism with these properties (BIC, IR, and WBB) that guarantees a *constant fraction* of the first-best GFT.

Benchmarks and Simple Mechanisms To better explain related work and our results, we first introduce some benchmarks and simple mechanisms. For simplicity, we assume F and G are both continuous distributions supported on a bounded interval $[0, 1]$ with positive densities. We show in Remark 5.2.2 why this is without loss of generality. We slightly overload the notations below and use them to denote both the mechanisms and the gains-from-trade they obtain.

- FB: the first best. It captures the optimal GFT if the agents are not strategic and there is a trade whenever $v \geq c$. Formally, $\mathsf{FB} = \int_0^1 \int_c^1 (v - c)\, \mathrm{d}F(v)\, \mathrm{d}G(c)$.

- SB: the second best. It captures the optimal GFT achieved by any Bayesian incentive-compatible (BIC), individually rational (IR) and weakly budget balanced (WBB) mechanism.

- FixedP: the fixed-price mechanism. The mechanism sets a fixed price p, and there is a trade if buyer's value is at least p and seller's cost is at most p. Formally, $\mathsf{FixedP} = \max_p \int_0^p \int_p^1 (v - c)\, \mathrm{d}F(v)\, \mathrm{d}G(c)$. The mechanism is dominant-strategy incentive-compatible (DSIC, stronger than BIC), IR and strongly budget balanced (SBB, stronger than WBB).

- $\mathsf{SellerP}$: the seller-pricing mechanism. The mechanism delegates the pricing

[1] A mechanism is weakly budget balanced if the total payment of the participants is ex-post non-negative, i.e., the payment from the buyer is always at least the revenue of the seller. Moreover, a mechanism is strongly budget balanced if the total payment of the participants is ex-post exactly 0, i.e., the payment from the buyer is always the revenue of the seller.

power to the seller, who in turn posts a price r_c to maximize her profit with knowledge of c and F. The buyer then decides whether to buy depending on whether $v \geq r_c$. Formally, $\mathsf{SellerP} = \int_0^1 \int_{r_c}^1 (v - c)\,\mathrm{d}F(v)\,\mathrm{d}G(c)$, where $r_c \in \arg\max_p (p-c)\cdot(1-F(p))$ is the price the seller sets when she has cost c. The mechanism is BIC, IR and SBB.

- BuyerP: the buyer-pricing mechanism. Symmetric to $\mathsf{SellerP}$, the mechanism delegates the pricing power to the buyer, who sets a price r'_v to maximize his utility, and then the seller decides whether to sell based on whether $c \leq r'_v$. Formally, $\mathsf{BuyerP} = \int_0^1 \int_0^{r'_v} (v-c)\,\mathrm{d}G(c)\,\mathrm{d}F(v)$, where $r'_v \in \arg\max_p (v-p)\cdot G(p)$. The mechanism is BIC, IR and SBB.

With these notations, the result of [MS83] demonstrates $\mathsf{SB} < \mathsf{FB}$ whenever the supports of F and G "overlap". It has remained an open question ever since, on how far apart SB and FB can be. Specifically, is it the case that SB is always at least a constant fraction of FB? We answer this question in the positive. In particular, the better of (or a randomization over) $\mathsf{SellerP}$ and BuyerP guarantees at least 10% of the first-best gains-from-trade.

Theorem 5.1.1. $\mathsf{FB} \leq 2\cdot\mathsf{SellerP} + 8\cdot\mathsf{BuyerP} \leq 10\cdot\max(\mathsf{SellerP},\mathsf{BuyerP}) \leq 10\cdot\mathsf{SB}$.

Remark 5.1.2. In Section A.4, we improve the constant to get a 8.23-approximation.

5.1.1 Related Work

There has been a large body of work towards answering whether $\mathsf{SB} = \Omega(1)\cdot\mathsf{FB}$. In particular, [McA08] shows $\mathsf{FixedP} \geq \frac{1}{2}\cdot\mathsf{FB}$ when F and G are i.i.d. and [BM16] show $\mathsf{SellerP} \geq \frac{1}{e}\cdot\mathsf{FB}$ when F satisfies the monotone-hazard-rate condition. As for the negative direction, [BM16] show that for any $\varepsilon > 0$, $\mathsf{SB} \leq \left(\frac{2}{e}+\varepsilon\right)\cdot\mathsf{FB}$ in some instance. [BD16] show that for any $\varepsilon > 0$, there is some instance where $\mathsf{FixedP} < \varepsilon\cdot\mathsf{FB}$ (and further $\mathsf{FixedP} < \varepsilon\cdot\mathsf{SB}$).

A closely related question of *welfare* approximation has also been extensively studied before. In this line of work, the welfare is defined to be GFT plus the term $\mathbf{E}_{c\sim G}[c]$, i.e., the welfare gains buyer's value when trade happens and it gains seller's cost when no trade happens. Maximizing the welfare is equivalent to maximizing the gains-from-trade, but providing a constant approximation to first-best gains-from-trade is much harder than doing the same for first-best welfare. For instance, not having any trade at all already gives a non-zero (and often good) approximation to welfare, but it gives a zero approximation to gains-from-trade. [BD16] show a $\left(1-\frac{1}{e}\right)$-approximation to the first-best welfare. [KPV22] improve the approximation ratio to $1-\frac{1}{e}+10^{-4}$.

[BCWZ17] present a 2-approximation against second-best gains-from-trade; in particular, they show $\mathsf{SB} \leq \mathsf{SellerP} + \mathsf{BuyerP}$. Our result is "stronger" in the sense that we show the better of $\mathsf{SellerP}$ and BuyerP is a constant-approximation of FB and not only SB. [DFL$^+$21] and [KPV22] consider settings without full knowledge of the distributions.

The *double auction* is a generalized version of bilateral trade, where there are multiple buyers and sellers in the market. There has been a significant amount of work on characterizing and approximating the efficient solutions in related settings; see e.g. [CdKLT16, CGdK$^+$17, BCWZ17, BCGZ18, BMPLZ19, BGG20, CGMZ21].

5.2 Proof of Efficiency Approximation

Recall that F and G are independent continuous CDFs supported on $[0,1]$ with positive densities (see Remark 5.2.2 for why everything apart from independence is without loss of generality). Let $\mu(x)$ be the median of $F_{|\geq x}$ (where $F_{|\geq x}(z)=0$ for $z<x$ and $F_{|\geq x}(z)=\frac{F(z)-F(x)}{1-F(x)}$ for $z\geq x$), i.e., $\mu(x)=F^{(-1)}\left(\frac{1+F(x)}{2}\right)$. Let $\mu^{(k)}(\cdot)$ be the composition of k functions of $\mu(\cdot)$, and $\mu^{(-k)}(\cdot)$ be its inverse (0 if it does not

exist). We first give immediate bounds for FB, $\mathsf{SellerP}$ and BuyerP:

$$\begin{aligned}\mathsf{FB} &= \int_0^1 \int_c^1 (v-c)\,\mathrm{d}F(v)\,\mathrm{d}G(c)\\ &\le 2\int_0^1 \int_{\mu(c)}^1 (v-c)\,\mathrm{d}F(v)\,\mathrm{d}G(c),\end{aligned}$$

where the second step holds since $\mu(x)$ is the median of $F_{|\ge x}$, and we preserve the better half.

For $\mathsf{SellerP}$, we have

$$\mathsf{SellerP} \ge \int_0^1 \int_{\mu(c)}^1 \Big(\mu(c)-c\Big)\,\mathrm{d}F(v)\,\mathrm{d}G(c).$$

The right-hand side (RHS) is the seller's expected profit when setting the price at $\mu(c)$ for each c, which is at most her optimal profit and thus at most the gains-from-trade in the profit-optimal seller-pricing mechanism (which is the left-hand side (LHS)).

Similarly, for BuyerP, we have

$$\begin{aligned}\mathsf{BuyerP} &\ge \int_{\mu^{(2)}(0)}^1 \int_0^{\mu^{(-2)}(v)} \Big(v-\mu^{(-2)}(v)\Big)\,\mathrm{d}G(c)\,\mathrm{d}F(v)\\ &= \int_0^1 \int_{\mu^{(2)}(c)}^1 \Big(v-\mu^{(-2)}(v)\Big)\,\mathrm{d}F(v)\,\mathrm{d}G(c).\end{aligned}$$

Again, the RHS is the buyer's expected utility when setting the price at $\mu^{(-2)}(v)$ for each v, which is at most the gains-from-trade in the utility-optimal buyer-pricing mechanism (the LHS).

For each cost c, define $\mathsf{FB}(c) = 2\int_{\mu(c)}^1 (v-c)\,\mathrm{d}F(v)$, $\mathsf{SellerP}(c) = \int_{\mu(c)}^1 \Big(\mu(c) - c\Big)\,\mathrm{d}F(v)$ and $\mathsf{BuyerP}(c) = \int_{\mu^{(2)}(c)}^1 \Big(v-\mu^{(-2)}(v)\Big)\,\mathrm{d}F(v)$. The bounds above simplify to:

$$\mathsf{FB} \le \int_0^1 \mathsf{FB}(c)\,\mathrm{d}G(c), \quad \mathsf{SellerP} \ge \int_0^1 \mathsf{SellerP}(c)\,\mathrm{d}G(c), \quad \mathsf{BuyerP} \ge \int_0^1 \mathsf{BuyerP}(c)\,\mathrm{d}G(c).$$

We now show that, for any c, $\mathsf{FB}(c) \leq 2 \cdot \mathsf{SellerP}(c) + 8 \cdot \mathsf{BuyerP}(c)$, thus proving our result. The crux of our proof is a partition of value space by quantile. We decompose $\mathsf{FB}(c)$ into two parts and bound them by $\mathsf{SellerP}(c)$ and $\mathsf{BuyerP}(c)$ separately.

Lemma 5.2.1. *For any c, $\mathsf{FB}(c) \leq 2 \cdot \mathsf{SellerP}(c) + 8 \cdot \mathsf{BuyerP}(c)$.*

Proof. We first have

$$\begin{aligned}
\mathsf{FB}(c) &= 2\int_{\mu(c)}^{1} (v - c)\,\mathrm{d}F(v) \\
&= 2\int_{\mu(c)}^{1} \Big(\mu(c) - c\Big)\,\mathrm{d}F(v) + 2\int_{\mu(c)}^{1} \Big(v - \mu(c)\Big)\,\mathrm{d}F(v) \\
&= 2 \cdot \mathsf{SellerP}(c) + 2\int_{\mu(c)}^{1} \Big(v - \mu(c)\Big)\,\mathrm{d}F(v).
\end{aligned}$$

We now proceed to show that $\int_{\mu(c)}^{1} \Big(v - \mu(c)\Big)\,\mathrm{d}F(v) \leq 4 \cdot \mathsf{BuyerP}(c)$. Observe that

$$\begin{aligned}
\int_{\mu(c)}^{1} \Big(v - \mu(c)\Big)\,\mathrm{d}F(v) &= \sum_{k=1}^{\infty} \int_{\mu^{(k)}(c)}^{\mu^{(k+1)}(c)} \Big(v - \mu(c)\Big)\,\mathrm{d}F(v) \\
&\leq \sum_{k=1}^{\infty} \int_{\mu^{(k)}(c)}^{\mu^{(k+1)}(c)} \Big(\mu^{(k+1)}(c) - \mu(c)\Big)\,\mathrm{d}F(v) \\
&= \sum_{k=1}^{\infty} \sum_{t=1}^{k} \int_{\mu^{(k)}(c)}^{\mu^{(k+1)}(c)} \Big(\mu^{(t+1)}(c) - \mu^{(t)}(c)\Big)\,\mathrm{d}F(v) \\
&= \sum_{t=1}^{\infty} \sum_{k=t}^{\infty} \int_{\mu^{(k)}(c)}^{\mu^{(k+1)}(c)} \Big(\mu^{(t+1)}(c) - \mu^{(t)}(c)\Big)\,\mathrm{d}F(v) \\
&= 2\sum_{t=1}^{\infty} \int_{\mu^{(t)}(c)}^{\mu^{(t+1)}(c)} \Big(\mu^{(t+1)}(c) - \mu^{(t)}(c)\Big)\,\mathrm{d}F(v).
\end{aligned}$$

The last step uses the definition of $\mu(\cdot)$: $k = t$ counts for half the probability of all $k \geq t$.

On the other hand, for $\mathsf{BuyerP}(c)$, recall that $\mathsf{BuyerP}(c) = \int_{\mu^{(2)}(c)}^{1} \left(v - \mu^{(-2)}(v)\right) \mathrm{d}F(v)$ and we have

$$\begin{aligned}
\int_{\mu^{(2)}(c)}^{1} \left(v - \mu^{(-2)}(v)\right) \mathrm{d}F(v) &= \sum_{t=2}^{\infty} \int_{\mu^{(t)}(c)}^{\mu^{(t+1)}(c)} \left(v - \mu^{(-2)}(v)\right) \mathrm{d}F(v) \\
&\geq \sum_{t=2}^{\infty} \int_{\mu^{(t)}(c)}^{\mu^{(t+1)}(c)} \left(\mu^{(t)}(c) - \mu^{(t-1)}(c)\right) \mathrm{d}F(v) \\
&= \sum_{t=1}^{\infty} \int_{\mu^{(t+1)}(c)}^{\mu^{(t+2)}(c)} \left(\mu^{(t+1)}(c) - \mu^{(t)}(c)\right) \mathrm{d}F(v) \\
&= \frac{1}{2} \sum_{t=1}^{\infty} \int_{\mu^{(t)}(c)}^{\mu^{(t+1)}(c)} \left(\mu^{(t+1)}(c) - \mu^{(t)}(c)\right) \mathrm{d}F(v),
\end{aligned}$$

where the last step again applies the definition of $\mu(\cdot)$. Therefore, we conclude the proof that $\mathsf{FB}(c) \leq 2 \cdot \mathsf{SellerP}(c) + 8 \cdot \mathsf{BuyerP}(c)$. $\square$

Remark 5.2.2. It is without loss of generality to consider bounded continuous distributions with positive densities: For any general distributions F and G on $[0,1]$, there is a sequence of continuous distributions on $[0,1]$ with positive densities converging to it in the Lévy metric. All of FB, seller's profit in SellerP and buyer's utility in BuyerP are continuous in F and G (in the Lévy metric). It is also without loss of generality to consider bounded supports: As long as FB is finite, almost all of the contribution comes from $(v, c) \in [-M, M]^2$ for some M.

Remark 5.2.3. By symmetry, we also have $\mathsf{FB} \leq 8 \cdot \mathsf{SellerP} + 2 \cdot \mathsf{BuyerP}$. Thus, using SellerP with probability α and BuyerP with probability $1 - \alpha$ for any $\alpha \in [0.2, 0.8]$ gives a 10-approximation to FB.

Remark 5.2.4. We can use a parameter to control the quantile of $\mu(x)$ (instead of using the median) in $F_{|\geq x}$. This improves the approximation constant to 8.23. The details are deferred to Section A.4.

6

Approximate Efficiency and Fairness in Auctions via Persuasion

6.1 Introduction

Consider the case where a seller is selling an item to a buyer, whose private value V is drawn from some known distribution $\mathcal{D}$. The seller has the pricing power, i.e., she can announce a price to the buyer, and then the buyer will decide whether he will buy it. The overall social welfare is maximized – i.e., the trade is efficient – when the seller sells the item for \$0, assuming the seller has no cost for the item. In contrast, to maximize the (average) revenue, the seller's optimal strategy is to sell at a revenue-maximizing price, which may lead to welfare loss due to the item going unsold.

More generally, in a single-item Bayesian auction with n buyers with independent private valuations, an efficient mechanism is the second-price (or VCG) auction, which always gives the item to the highest-valued buyer. In contrast, even when the buyers have i.i.d. regular valuations, the revenue-optimal mechanism was shown by [Mye81] to be a second-price auction *with a reserve price*; this may lead to the item going unsold. The situation is more complex for non-regular distributions, and/or non-i.i.d. buyers, where the revenue-optimal mechanism may in addition sell the item to a buyer with lower value than the highest, leading to additional welfare loss. We visualize this via a *revenue-CS trade-off* diagram (Fig. 6.1b), where, for different mechanisms and value distributions, we plot expected consumer-surplus (i.e., value minus payment), denoted CS, versus expected seller-revenue, denoted by $\mathcal{R}$. Any welfare-maximizing mechanism including VCG (point V) lies on the $\mathcal{R} + \text{CS} = \mathcal{W}^*$ line. In contrast, Myerson's mechanism (point M) has revenue $\mathcal{R}^M$ greater than that under VCG, but can also lie below the maximum-welfare line.

Information Intermediary. Now consider the same setting, but with an additional *information intermediary*: a third-party who knows the *true* buyer values $\vec{V} = (V_1, V_2, \dots, V_n)$ and can provide a "signal" or side-information to the seller and the buyers. Both the signal and signaling scheme are common knowledge to all agents (buyers and seller), who can thus use Bayes' rule to update the prior over valuations given the signal. The signal "re-shapes" the joint prior over the buyer valuations in a *Bayes-plausible* manner (i.e., such that the posterior averaged over signals equals the prior). Though the intermediary can modulate information, it does not control the mechanism, which still resides with the seller. Such a setting is motivated by ad exchanges, where the platform (or intermediary) acts only as a clearinghouse, and does not itself run a mechanism. Therefore, given the signal, the seller then proposes the revenue-maximizing mechanism, and buyers bid optimally, under the posterior distribution. We illustrate this in Fig. 6.1a.

Formally, consider a setting where n buyers have independent private valuations $\vec{V}$ drawn from a distribution $\mathcal{D} = \mathcal{D}_1 \times \mathcal{D}_2 \times \cdots \times \mathcal{D}_n$. The valuations $\vec{V}$ are known to the intermediary, who maps them to a signal σ via a public signaling scheme $\mathcal{Z}$. (For our positive results, only the seller needs to receive the signal. It does not matter if the buyers receive them.) Given σ, all agents compute the posterior $\mathcal{S}$ over buyer values; note these can now be correlated. The seller then proposes a mechanism $\mathcal{M}_{\mathcal{S}}$ (comprising allocation and payment rules) which maximizes its expected revenue assuming buyers act in a manner which is ex-post incentive-compatible (IC) and interim individually-rational (IR) given $\mathcal{S}$. If σ is such that $\mathcal{S} = \mathcal{D}$, then $\mathcal{M}_{\mathcal{S}}$ is Myerson's auction (point M in Fig. 6.1b); on the other hand, if the signal fully reveals $\vec{V}$, then the seller can extract full surplus (i.e., get revenue $\mathcal{W}^*$, point A in Fig. 6.1b). Moreover, the seller gets revenue at least $\mathcal{R}^M$ under any signaling scheme, as she can always ignore the signal (see Section 6.2). Thus any signaling scheme $\mathcal{Z}$ must give a point in the shaded triangle with consumer surplus $\text{CS}(\mathcal{Z})$ and revenue $\mathcal{R}(\mathcal{Z})$, and the maximum possible surplus $\textsc{Opt}$ is achieved at point O in Fig. 6.1b. Now we can ask:

> *What revenue-CS trade-offs can an information intermediary achieve via signaling? More specifically, what is the maximum possible consumer surplus that is achievable?*

In the single-buyer case, the seminal work of [BBM15] completely answer these questions by showing that the entire shaded region is always achievable. In particular, the point O is met by a simple signaling scheme where the revenue is *exactly* $\mathcal{R}^M$, and the item is *always* sold thus the mechanism is efficient.

In this chapter we study the effectiveness of an information intermediary in a multi-buyer (i.e., $n \geq 2$ buyers) Bayesian auction. In brief, we expose a sharp separation between the single and multi-buyer settings, as in the latter, no signaling scheme can guarantee more than a constant fraction of the optimal consumer surplus ($\textsc{Opt}$ in Fig. 6.1b).

On the positive side, we obtain a novel yet simple signaling scheme with strong approximation guarantees for a wide range of settings.

While our main focus is on theoretical results, our work has broader practical relevance. Consider an agency like the FCC with privileged information about bidders in a spectrum auction, or a bid optimizer working for multiple competing clients in an ad-exchange. These intermediaries have private information about the buyers, and can selectively release it to influence the auction. For instance, in an ad exchange, the platform running the exchange (intermediary) typically uses machine learning and advertiser features to infer true valuations. However, the pricing rules are decided by the publishers (seller) and not the exchange. For its own long term viability, the platform clearly has incentives to make both parties – publisher and advertisers – as happy as possible, and would therefore like to release information selectively to the publisher in order to maximize advertiser happiness (consumer surplus) while keeping publisher happiness (revenue) at least what it is without its presence. In effect, we use the alternate view of the market segmentation problem in [BBM15] as a special case of a signaling problem where a more informed intermediary works for the benefit of the buyers.

Our work also fits in a broader space of multi-criteria optimization where a third-party platform or government agency can release information about agents to a principal in charge of an activity such as admissions or hiring, so as to trade-off the principal's objective such as maximizing quality of hire (or revenue in our setting), with a societal objective such as fairness or diversity (or the consumer surplus in our setting).

6.1.1 Our Results

We consider a single-item auction with n buyers with discrete valuations. We assume the buyer valuations are independent, so $\mathcal{D} = \mathcal{D}_1 \times \mathcal{D}_2 \times \cdots \times \mathcal{D}_n$, where $\mathcal{D}_i$ has support size $\mathcal{K}_i$, and the size of the union of the supports is $\mathcal{K}$.

Our first set of results (Section 4.4) shows a sharp demarcation between the cases

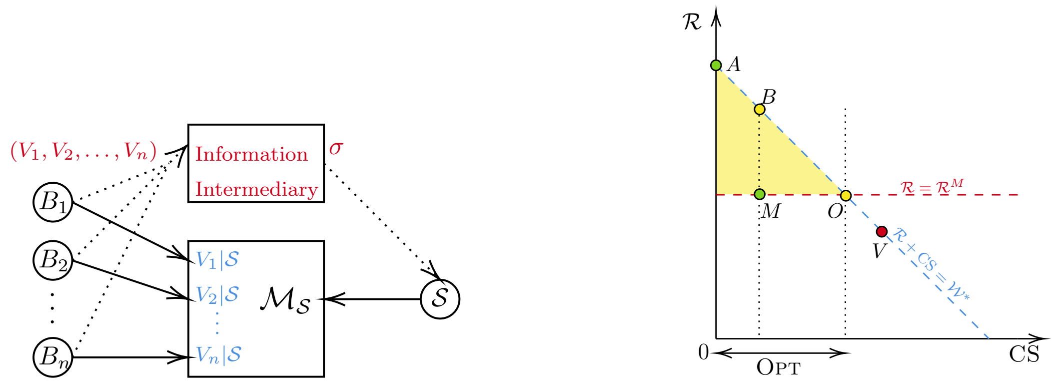

(a) Auction with Information Intermediary

(b) Revenue-CS Trade-off Diagram

Figure 6.1: (a) The auction with information intermediary setting, where the intermediary has full knowledge of valuations $\vec{V}$, and can use this to provide a signal σ to the seller and the buyers. The seller then uses the revenue-optimal mechanism $\mathcal{M}_{\mathcal{S}}$ for the posterior distribution over valuations $\mathcal{S}$ given σ.
(b) Two-dimensional space of seller revenue, $\mathcal{R}$, and consumer surplus, CS, of different signaling schemes. The points M and V correspond to Myerson's and VCG mechanisms, and the point A corresponds to selling to the highest-valued buyer at her value when the seller has full information.

of $n = 1$ and $n \geq 2$ buyers. In contrast with the former case (where signaling achieves the entire shaded region in Fig. 6.1b), we show in the latter case, the entire segment BO is not achievable; indeed, the only achievable points on segment AO are arbitrarily close to A. Therefore, achieving full welfare requires sacrificing an arbitrarily large fraction of consumer surplus compared to the no-signaling baseline.

Theorem 6.1.1 (Proved in Section A.5)**.** *For any given constant $\varepsilon > 0$, there are instances with $n = 2$ buyers each with $\mathcal{K}_i = 2$, where any signaling scheme $\mathcal{Z}$ under which the revenue-optimal auction obtains full welfare (i.e., allocates to highest-value buyer), has $\mathrm{CS}(\mathcal{Z}) \leq \varepsilon \cdot \mathrm{CS}(\mathcal{D})$, where $\mathrm{CS}(\mathcal{D})$ is the consumer surplus of Myerson's auction without signaling.*

We next ask if we can sacrifice on welfare, but raise a consumer surplus arbitrarily close to Opt? We again answer in the negative, and show a lower bound of 2 on the approximation ratio.

Theorem 6.1.2 (Proved in Section A.5)**.** *For any constant $\varepsilon > 0$, there are problem instances with $n = 2$ buyers each with $\mathcal{K}_i = 2$, where any signaling scheme $\mathcal{Z}$ has* $\mathrm{CS}(\mathcal{Z}) \leq \left(\frac{1}{2} + \varepsilon\right) \cdot \textsc{Opt}$.

We note that the above results are *existential impossibility results*, and do not depend on the complexity of the signaling scheme.[1] Overall, the negative results in Theorems 6.1.1 and 6.1.2 strongly suggest that in this setting, the focus should be on approximating the consumer surplus.

The situation improves in Section 6.2.3 when we restrict to $\mathcal{D}_i$ that are identical and (discrete-)regular. Here, we first circumvent Theorem 6.1.1 by showing a simple signaling scheme that achieves the point B (i.e., optimal welfare, and same consumer surplus as under Myerson's auction).

One problem that remains, however, is that Myerson's auction may have arbitrarily poor CS: for example, if the $\mathcal{D}_i$'s are regular, and chosen such that the reserve price is the highest value in the support, then $\mathrm{CS} = 0$, while $\textsc{Opt} > 0$ (and so the approximation factor of Myerson's auction is unbounded). Indeed, even restricting to MHR priors, one can construct instances where the reserve price is close to the maximum value in the support, leading Myerson's auction to have vanishing CS relative to $\textsc{Opt}$. This is one reason why getting any non-trivial approximation to $\textsc{Opt}$ is challenging, and we present more discussion in Section 6.3.1.

In Section 6.3, we present our main technical result, where we show that when buyers' valuations are drawn from i.i.d. regular distributions, then a simple signaling scheme achieves a constant-approximation to $\textsc{Opt}$. In more detail, our $\textsc{Rank}_t$ signaling scheme is based on two simple but critical steps: First, the intermediary can use its knowledge of agent valuations to perform a *pre-screening* step that eliminates all but the top-t buyers (for a carefully chosen t). Second, given the top t buyers, it

[1] Our proofs also imply the same lower bounds when the seller is constrained to run an *ex-post* IR mechanism, provided the intermediary's signals induce a product-form posterior distribution.

can then choose a uniform buyer among this set to serve as a *hold-out buyer*, who the seller can sell to in case she is unable to raise sufficient revenue from the remaining $t-1$ buyers via an auction; this can be achieved by using the single-buyer signaling scheme of [BBM17] on the chosen buyer. Using a combination of these two ideas, we get the following:

Theorem 6.1.3 (Proved in Section 6.3). *When the $\mathcal{D}_i$'s are identical and regular, there is a signaling scheme that achieves an $O(1)$-approximation to the optimal consumer surplus* OPT*, and has computation time polynomial in n and $\mathcal{K}$.*

6.1.2 Intuition and Techniques

For any n, the optimal signaling scheme for maximizing surplus can be obtained via an infinite-sized linear program (see Eq. (A.5) in Section 4.4) with variables for every possible signal, i.e., every possible joint distribution over buyer valuations. Further, for each such signal, the quantity of interest is the consumer surplus of the revenue-optimal auction given the signal. For $n = 1$ case, Bergemann et al. show this LP has a special structure in that it admits a basis comprising of "equal-revenue distributions" containing the revenue-maximizing price (see Section 6.2.2). Our work shows that this breaks down for optimal auctions with signaling involving $n \geq 2$ buyers.

To understand why things change dramatically from $n = 1$ to $n \geq 2$ buyers, in the former case, the optimal mechanism is a simple posted price scheme and its revenue is continuous in the distribution $\mathcal{D}$. However, with multiple buyers, the optimal auction does not have simple structure even for independent buyers (see Algorithm 4), and we need to analyze the consumer surplus of this auction, which can be a discontinuous function of the prior. (See Section 4.4 for examples.) Further, for correlated buyers, the revenue of the auction itself may not be continuous in the prior! Indeed, a celebrated result of [CM88] shows that slightly perturbing an independent prior to a correlated one can discontinuously increase the revenue to $\mathcal{W}^*$, hence decreasing consumer surplus to 0. (See Theorem 6.2.1 in Section 6.2.) This makes it tricky to reason about the optimal signaling scheme, leading to the gap between our upper and lower bounds.

In more detail: Our proofs of Theorem 6.1.1 and Theorem 6.1.2 use a special case of the Crémer-McLean characterization [CM88]: for $n = 2$ buyers each with $\mathcal{K}_i = 2$, under any non-independent prior the seller can extract full social surplus as revenue. This lets us focus on signaling schemes where buyers' posteriors given each signal are product distributions. Using Myerson's characterization of the optimal auction for discrete valuations [Elk07], we show a structural characterization that reduces the space of optimal signals to a sufficiently simple form, yielding the desired counterexamples. Note that we still need to reason about a large space of possible product distributions as signals, which makes our constructions quite non-trivial.

The technically most interesting result is the $O(1)$-approximation signaling scheme for i.i.d. buyers (Theorem 6.1.3 in Section 6.3). The challenge is the following: Even if we restrict the space of signals so that the posteriors are (non-identical) product distributions, this space is still infinite size, with CS being a discontinuous function in this space. Our signaling scheme in Section 6.3 balances the trade-off between revealing enough information about valuations so that the item is sold to a high-value buyer, and revealing too much information such that the seller extracts all the surplus. Balancing these is delicate; nevertheless, our final scheme is simple with polynomial computation time and signal complexity. We present more intuition in Section 6.3.1, where we argue that the guarantee in Theorem 6.1.3 cannot be achieved in a straightforward fashion.

6.1.3 Related Work

The general problem of information structure design considers how sharing additional information can influence the outcome of a mechanism. Different variants of this problem have been formulated and studied; we refer the reader to [BM19, Dug17] for surveys. Of particular importance to us is the *Bayesian persuasion problem* formulated by [KG11], where a receiver selects a utility-maximizing action based on incomplete information about the state of nature. A sender who knows the state of nature can signal side-information to the receiver so that the action taken by the receiver is utility-maximizing for the sender. This general problem has been widely studied in different domains such as monopoly pricing and advertising [BBM15, CH14, XRDT15, HS22]. For this problem, there is a distinction between existence and computational results, and the work of [DX19] studies the computational complexity of finding the optimal signaling scheme under different input models.

The restriction of our problem to one buyer is the monopoly pricing problem. Here, the intermediary is the sender whose utility is consumer surplus, and the seller is the receiver whose action space is take-it-or-leave-it prices and whose utility is revenue. Beginning with the work of [BBM15], several works [DKQ16, STZ18, CDHW20, CEF$^+$20, HS22, KM22, MPLW22] have considered various extensions and modifications to this basic problem. Unlike monopoly pricing where the buyer is perfectly informed, in our setting, not only the seller, but also all the buyers are receivers, in the sense that they have imperfect knowledge of the true valuations of other buyers, and modify their respective bidding strategies in response to the intermediary's signal to maximize their own utilities. Our setting is therefore a Bayesian persuasion problem with multiple receivers, and this aspect makes it significantly more complex.

There has been work on signaling in auctions that cannot be modeled as Bayesian persuasion, i.e., in which the common signal is not generated by an intermediary who knows all the true values of the buyers. For instance, in the work of [BP07], the auctioneer has perfect information about buyer valuations and controls the precision to which buyers can learn it, and in the work of [FLLT18], the seller's signal is drawn from a distribution that is correlated with the buyer's value, In both these works, the goal is to maximize seller revenue. Finally, [STZ19] studies equilibria of optimal auctions when each buyer commits to a signaling scheme with imperfect knowledge of other buyers' valuations, while [BBM17] studies equilibria in first price auctions when buyers are provided correlated signals about other buyers' valuations. In contrast with the former, our work considers a richer space of signals via an information intermediary, while compared to the latter, in our setting the seller's mechanism is not fixed, but is instead also a function of the information structure.

6.2 Preliminaries

We consider Bayesian single-item auctions with n buyers, with independent private valuations $\vec{V} = (V_1, V_2, \ldots, V_n)$ drawn from a known product distribution $\mathcal{D} = \mathcal{D}_1 \times \cdots \times \mathcal{D}_n$. Unless otherwise stated, we present our results for the setting in which each $\mathcal{D}_i$ is discrete. We denote by $\mathcal{K}_i$ the size of the support of $\mathcal{D}_i$, and by $\mathcal{K}$ the size of the union of these supports.

For distribution $\mathcal{D}_i$, we use f_{D_i} to denote its probability mass function, and define $S_{D_i}(x) = \Pr_{V_i \sim D_i}[V_i \geq x]$ and $F_{D_i}(x) = \Pr_{V_i \sim D_i}[V_i \leq x]$. For a joint distribution $\mathcal{D}$ and vector $\vec{v}$, we use $\Pr[\mathcal{D} = \vec{v}] = f_{\mathcal{D}}(\vec{v})$ as shorthand for denoting the probability of $\vec{v}$ drawn from $\mathcal{D}$.

6.2.1 Revenue-Maximizing Auctions

Given any shared prior $\mathcal{D}'$ on the valuations of the buyers, which in the case of signaling, can be different from $\mathcal{D}$ and arbitrarily correlated, the seller runs an optimal (revenue maximizing) auction that satisfies ex-post incentive compatibility and *interim* individual rationality. Let $\mathcal{D}'_i$ denote the marginal distribution of buyer i. Using the revelation principle [Mye81], the optimal auction is specified by an allocation rule $x^*(\vec{v}) \geq 0$ and a payment rule $\theta^*(\vec{v})$ (that can be positive or negative) given any realized valuation profile $\vec{v}$. The quantity $x^*_i(\vec{v})$ is the probability buyer i gets the item given the valuation profile $\vec{v}$. The allocation is always non-negative, while the price could be negative. The auction maximizes revenue:

$$\mathcal{R}(\mathcal{D}') = \max_{\vec{x},\vec{p}} \sum_{\vec{v}} \Pr[\mathcal{D}' = \vec{v}] \cdot \sum_i \theta_i(\vec{v})$$

subject to the following constraints. The *ex-post* incentive compatibility (IC) constraint states that the utility of any buyer does not increase by misreporting its valuation. Formally, for every buyer i, for every value q, and for every valuation vector $\vec{v} = (q, \vec{v}_{-i})$ (where $\vec{v}_{-i}$ denotes the valuations of the other buyers), and every other possible report r,

$$q \cdot x_i(q, \vec{v}_{-i}) - \theta_i(q, \vec{v}_{-i}) \geq q \cdot x_i(r, \vec{v}_{-i}) - \theta_i(r, \vec{v}_{-i}).$$

The *interim* individual rationality (IR) constraint says that for any buyer i and any value q, the expected utility under the mechanism is non-negative:

$$\sum_{\vec{v}_{-i}} \frac{\Pr[\mathcal{D}' = (q, \vec{v}_{-i})]}{\Pr[\mathcal{D}'_i = q]} \cdot (q \cdot x_i(q, \vec{v}_{-i}) - \theta_i(q, \vec{v}_{-i})) \geq 0.$$

Finally, we have the constraint that the item is allocated probabilistically to at most one buyer:

$$\sum_i x_i(\vec{v}) \leq 1, \qquad \forall \vec{v}.$$

For any prior $\mathcal{D}'$, let $(\mathcal{R}(\mathcal{D}'), \mathcal{W}(\mathcal{D}'), \mathrm{CS}(\mathcal{D}'))$ denote the expected revenue, welfare (or *total surplus*) and consumer surplus under the revenue-maximizing auction.

Then we have $\mathrm{CS}(\mathcal{D}') = \mathcal{W}(\mathcal{D}') - \mathcal{R}(\mathcal{D}')$, and:

$$\mathcal{R}(\mathcal{D}') = \sum_{\vec{v}} \Pr[\mathcal{D}' = \vec{v}] \cdot \sum_i \theta_i^*(\vec{v}) \qquad \text{and} \qquad \mathcal{W}(\mathcal{D}') = \sum_{\vec{v}} \Pr[\mathcal{D}' = \vec{v}] \cdot \sum_i v_i x_i^*(\vec{v}),$$

where $x^*(\vec{v}) \geq 0$ and $\theta^*(\vec{v})$ are the allocation rule and the payment rule of the optimal auction given any realized valuation profile $\vec{v}$.

Our work builds on two special cases – independent valuations, and full surplus extraction.

Optimal auction for independent valuations. When $\mathcal{D} = \mathcal{D}_1 \times \cdots \times \mathcal{D}_n$ is a product distribution, the optimal auction has a simple form given by [Mye81]. For distribution $\mathcal{D}_i$ with support $z_1 < z_2 < \cdots < z_k$, its *virtual value* function $\varphi_{\mathcal{D}_i}$ is defined as:

$$\varphi_{\mathcal{D}_i}(z_k) = z_k \qquad \text{and} \qquad \varphi_{\mathcal{D}_i}(z_\ell) = z_\ell - (z_{\ell+1} - z_\ell)\frac{S_{\mathcal{D}_i}(z_{\ell+1})}{f_{\mathcal{D}_i}(z_\ell)}, \qquad \forall \ell < k. \quad (6.1)$$

If buyer i is the only buyer in the system, the optimal auction sets a fixed price, and the buyer buys the item when her valuation is at least this price. The *reserve price* of $\mathcal{D}_i$, denoted $r_{\mathcal{D}_i}$ is the smallest value r in the support of $\mathcal{D}_i$ that maximizes the corresponding revenue $rS_{\mathcal{D}_i}(r)$. It is easy to check that $\varphi_{\mathcal{D}_i}(r_{\mathcal{D}_i}) \geq 0$.

Throughout this paper, we assume the distributions $\mathcal{D}_i$ are *regular*, so that $\varphi_{\mathcal{D}_i}(z)$ is a non-decreasing function of z. Therefore, for all $v < r_{\mathcal{D}_i}$, we have $\varphi_{\mathcal{D}_i}(v) < 0$.

For discrete regular distributions, Myerson's auction takes the form [Elk07] in Algorithm 4. Note that this auction is also *ex-post* IC and IR.

Algorithm 4 Myerson's Auction with prior $\mathcal{D}$ and valuations $\vec{v}$.

1: Sort the buyers in decreasing order of $q_i = \varphi_{\mathcal{D}_i}(v_i)$. Assume no two values are identical (can be ensured by using a fixed tie-breaking rule).
2: Allocate to the bidder j with highest virtual value q_j, provided $q_j \geq 0$.
3: Let m be the bidder with second highest virtual value, and let $w = \max(0, q_m)$.
4: Charge j the smallest value z in the support of $\mathcal{D}_j$ such that $\varphi_{\mathcal{D}_j}(z) > w$.

Extracting full surplus as revenue. At the other extreme, a celebrated result of [CM88] shows that for distributions $\mathcal{D}'$ which are "sufficiently correlated", the optimal auction extracts full surplus (i.e., the revenue equals the maximum valuation in each valuation profile). Formally, the result requires that for each agent, their conditional distribution over others' values given their own value is full rank; for our purposes, we require a restriction of their result to $n = 2$ buyers, each with two possible valuations.

Theorem 6.2.1 ([CM88])**.** *For $n = 2$ buyers, where each buyer i has $\mathcal{K}_i = 2$ and the joint distribution over the valuations is $\mathcal{D}'$, the seller (who faces an interim IR constraint) can extract the entire social welfare (i.e. get expected revenue equal to the expected value of the maximum of the buyer's valuations) when $\mathcal{D}'$ is a correlated (i.e. not independent) distribution.*

6.2.2 Auctions with an Information Intermediary

We next formalize the model of an information intermediary illustrated in Fig. 6.1a. Since the effect of the intermediary's signal is captured by the resulting posterior distribution over valuations, for ease of notation, we henceforth use "signal" to refer to a distribution $\mathcal{S}$ over valuations.

A signaling scheme $\mathcal{Z} = \{\gamma_q, \mathcal{S}_q\}_{q\in[m]}$ comprises a collection of *signals* (i.e., joint distributions over valuations) $\mathcal{S}_1, \mathcal{S}_2, \ldots, \mathcal{S}_m$ and corresponding non-negative weights $\gamma_1, \gamma_2, \ldots, \gamma_m$. The scheme $\mathcal{Z}$ is feasible (or "Bayes plausible" [KG11]) if it satisfies $\sum_q \gamma_q = 1$ and $\sum_q \gamma_q \mathcal{S}_q = \mathcal{D}$. The intermediary commits to scheme $\mathcal{Z}$ before the auction, and it is known to the seller and all buyers.

The intermediary maps observed valuation profile $\vec{v} \sim \mathcal{D}$ to signal $\mathcal{S}_q$ with probability $\frac{\gamma_q \Pr[\mathcal{S}_q=\vec{v}]}{\Pr[\mathcal{D}=\vec{v}]}$. The seller uses $\mathcal{S}_q$ as the shared prior and runs an optimal auction on the buyers. Note that though $\mathcal{D}$ is a product distribution, the $\{\mathcal{S}_q\}$ can be correlated. Abusing the notations introduced earlier in Section 6.2.1, we denote the revenue generated by signaling scheme $\mathcal{Z}$ as $\mathcal{R}(\mathcal{Z}) = \sum_q \gamma_q \mathcal{R}(\mathcal{S}_q)$, its consumer surplus by $\mathrm{CS}(\mathcal{Z}) = \sum_q \gamma_q \mathrm{CS}(\mathcal{S}_q)$, and its welfare by $\mathcal{W}(\mathcal{Z}) = \sum_q \gamma_q \mathcal{W}(\mathcal{S}_q)$.

When $\mathcal{D}$ is a product distribution, the revenue from any signaling scheme must be at least the optimal revenue of Myerson's auction without signaling, $\mathcal{R}(\mathcal{D})$. To see this, we note that Myerson's auction on $\mathcal{D}$ is ex-post IC and IR. This means that this allocation and payment rule is still a feasible (interim IC and IR) mechanism conditioned on receiving any signal, completing the argument. Therefore, the consumer surplus $\mathrm{CS}(\mathcal{Z})$ under any signaling scheme $\mathcal{Z}$ is bounded by the difference of the maximum possible welfare $\mathcal{W}^* = \mathbf{E}_{\vec{V}\sim\mathcal{D}}[\max_i V_i]$ and the maximum revenue without signaling $\mathcal{R}(\mathcal{D})$. We henceforth denote this bound as $\textsc{Opt}$, which is defined as follows:

$$\textsc{Opt} = \mathcal{W}^* - \mathcal{R}(\mathcal{D}).$$

We say that $\mathcal{Z}$ is a τ-approximation signaling scheme if $\mathrm{CS}(\mathcal{Z}) \geq \frac{\textsc{Opt}}{\tau}$. Our goal is to find the best approximation factor τ via a signaling scheme whose computation time is polynomial in n and $\mathcal{K}$. In the rest of the paper, we omit the dependence on $\mathcal{D}$ when clear from context.

Optimal signaling for a single buyer. For $n = 1$ buyer, [BBM15] present a signaling scheme with consumer surplus exactly equal to $\textsc{Opt}$ (i.e., implementing the point O in Fig. 6.1b. Their signaling scheme constructs distributions (signals) $\mathcal{S}_1, \mathcal{S}_2, \ldots, \mathcal{S}_m$ and assigns weights $\gamma_1, \gamma_2, \ldots, \gamma_m$ to them such that $\sum_q \gamma_q \mathcal{S}_q = \mathcal{D}$.

Let prior $\mathcal{D}$ takes value v_i with probability η_i, where $0 < v_1 < \cdots < v_k$. Let $\vec{\eta} = (\eta_1, \eta_2, \cdots, \eta_k)$. In each iteration ℓ, the algorithm constructs an *equal revenue* distribution $\mathcal{S}_\ell$ and subtracts it from the prior $\mathcal{D}$. This equal revenue distribution assigns positive probability $\eta_{i\ell}$ to v_i if $\eta_i > 0$ and assigns $\eta_{i\ell} = 0$ if $\eta_i = 0$. In $\mathcal{S}_\ell$, the seller raises equal revenue by setting the price to be any of the values v_i with $\eta_i > 0$. It is easy to see that the equal revenue condition specifies a unique distribution $\mathcal{S}_\ell$.

Note that since this signal is equal revenue, (we may assume) the seller sets the lowest value as price, so that the item always sells and the consumer surplus is maximum possible.

Let $\vec{\eta_\ell}$ be the probability vector of $\mathcal{S}_\ell$. We set the largest weight γ_ℓ such that $\vec{\eta} - \gamma_\ell \vec{\eta_\ell} \geq 0$. We update $\mathcal{D}$ by setting $\vec{\eta}$ to $\vec{\eta} - \gamma_\ell \vec{\eta_\ell}$, and increase ℓ by one. We repeat this till the support of $\mathcal{D}$ becomes empty. The $\{\gamma_\ell, \mathcal{S}_\ell\}$ specifies the signaling scheme. We illustrate this procedure by an example.

Example 6.2.2. Suppose the type space is $\{1,2,3\}$ and $\mathcal{D} = \langle \frac{1}{3}, \frac{1}{3}, \frac{1}{3} \rangle$ are the probabilities of these types. The monopoly price is $\theta = 2$ with revenue $\mathcal{R}(\mathcal{D}) = \frac{4}{3}$, while the point A in Fig. 6.1b has social welfare $\mathcal{R}(A) = \mathcal{W}^* = \mathbf{E}[\mathcal{D}] = 2$. Suppose $\mathcal{S}_1 = \langle \frac{1}{2}, \frac{1}{6}, \frac{1}{3} \rangle$ with $\gamma_1 = \frac{2}{3}$; $\mathcal{S}_2 = \langle 0, \frac{1}{3}, \frac{2}{3} \rangle$ with $\gamma_2 = \frac{1}{6}$; and $\mathcal{S}_3 = \langle 0, 1, 0 \rangle$ with $\gamma_3 = \frac{1}{6}$. It is easy to check that the monopoly price for each signal is the lowest price in its support so that the item always sells, and $\sum \gamma_i \mathcal{R}(\mathcal{S}_i) = \frac{4}{3}$. Therefore, $\sum \gamma_i \mathrm{CS}(\mathcal{S}_i) = 2 - \frac{4}{3} = \frac{2}{3} = \textsc{Opt}$, which corresponds to point O in Fig. 6.1b.

We henceforth use $\mathrm{BBM}(v, D)$ to refer to this scheme when the buyer has valuation distribution D and the realized value is $v \sim D$. Below we state some critical properties of the BBM scheme which we use in our results.

Lemma 6.2.3 (Implicit in [BBM15])**.** *For a single buyer with value distribution $\mathcal{D}$ (with reserve price $r_{\mathcal{D}}$), the* BBM *mechanism satisfies the following properties:*

1. *For any signal $\mathcal{S}_q$, $\varphi_{\mathcal{S}_q}(v) \geq 0$ for all v in the support of $\mathcal{S}_q$.*

2. $\mathrm{CS}(\mathrm{BBM}) = \textsc{Opt} \geq \Pr_{V\sim\mathcal{D}}[V < r_{\mathcal{D}}] \cdot \mathbf{E}_{V\sim\mathcal{D}}[V \mid V < r_{\mathcal{D}}] = \sum_{v<r_{\mathcal{D}}} v f_{\mathcal{D}}(v)$.

6.2.3 Achieving the Pareto-Frontier in the I.I.D. Case

We now ask if there are cases where we can circumvent Theorem 6.1.1 and maximize welfare while ensuring at least as much surplus as Myerson's auction (that is, achieve a point on the line BO in Fig. 6.1b). Note that Theorem 6.1.1 rules this out for non-i.i.d. distributions. Surprisingly, however, for i.i.d. regular distributions $\mathcal{D}_i$, the following simple signaling scheme turns out to be sufficient for achieving point B in Fig. 6.1b. Morally this shows why our lower bound constructions are delicate.

Suppose the common reserve price of $\mathcal{D}_i$ is r, and the maximum value of any buyer is v_m.

1. If $v_m < r$, then the intermediary reveals v_m and the identity of the highest buyer to the seller, who then sells to this bidder at price v_m.

2. If $v_m \geq r$, the intermediary only reveals the information that some buyer has value $\geq r$ (but does not reveal either v_m or the identity of the highest bidder). In this case, though the posterior is not a product distribution anymore, it can be shown that the seller's optimal auction remains the second price auction with reserve r.[2]

It is easy to check that the item always sells to the highest buyer, and the CS is exactly the same as in Myerson's auction without signaling, thereby achieving point B. Note however that this scheme does not give any guarantees on approximating CS itself, since the surplus of Myerson's auction is not an approximation to Opt. The question of approximating Opt is much more challenging even for the i.i.d. regular setting, and this is what we focus on in the next section.

6.3 Approximating Consumer Surplus: The i.i.d. Case

In this section, we present our main result (Theorem 6.1.3): An $O(1)$-approximation to Opt when the buyers' valuation distributions $\mathcal{D}_i$ are identical and regular.

[2] Note this is the only case when the seller gets non-zero revenue in the optimal auction for the original product distribution. Suppose for the purpose of contradiction that the seller can do better for this posterior, she can also do better than the optimal auction for the original distribution.

Recall we start with a product distribution $\mathcal{D} = \mathcal{D}_1 \times \mathcal{D}_2 \times \cdots \times \mathcal{D}_n$. Let $x_1 < x_2 < \cdots < x_{\mathcal{K}}$ be the supports of value distributions. $\mathcal{R}(\mathcal{D})$ denotes the revenue of the optimal auction (Algorithm 4) on $\mathcal{D}$, and $\textsc{Opt} = \mathbf{E}_{\vec{v}\sim\mathcal{D}}[\max_i v_i] - \mathcal{R}(\mathcal{D})$.

When $\mathcal{D}_i$'s are identical and regular, since the highest virtual-value and highest value buyers coincide, the optimal auction (Algorithm 4) assigns the item to the buyer with highest value if this value is above the common reserve price. Therefore, we decompose $\textsc{Opt}$ into two components:

- Myerson's surplus: The CS generated by Myerson's auction, denoted by $\mathrm{CS}(\mathcal{D})$.
- Non-allocation surplus: The loss in CS due to Myerson's auction not allocating the item, denoted by CS_0.

We therefore have $\textsc{Opt} = \mathrm{CS}(\mathcal{D}) + \mathrm{CS}_0$. In the remainder, we will present an $O(1)$-approximation for the non-allocation surplus CS_0, which will imply Theorem 6.1.3 when $\mathrm{CS}_0 \geq \mathrm{CS}(\mathcal{D})$. (When $\mathrm{CS}(\mathcal{D}) > \mathrm{CS}_0$, sending no signal is already a 2-approximation.)

6.3.1 Preliminaries and Intuition

Our approximation bound for the non-allocation surplus will also apply when $\mathcal{D}_i$ are independent but not necessarily identical or regular. Therefore, in the sequel, we will proceed assuming the more general case that $\mathcal{D}_i$'s are not necessarily identical or regular, and derive signaling schemes that approximate the non-allocation surplus CS_0 for this case.

We denote a realization from $\mathcal{D}$ by $\vec{v} = \{v_i\}$. Let $p_i = \Pr_{v_i\sim\mathcal{D}_i}[v_i < r_{\mathcal{D}_i}]$ for any buyer i, where $r_{\mathcal{D}_i}$ is the reserve price of $\mathcal{D}_i$. Let $Y_i = \mathcal{D}_{i|<r_{\mathcal{D}_i}}$ denote the distribution of $\mathcal{D}_i$ conditioned on being smaller than $r_{\mathcal{D}_i}$. Suppose we draw a sample independently from each distribution Y_i. Let Z_ℓ denote the distribution for the ℓ^{th} largest value among these n draws.

We first derive an expression for CS_0.

Lemma 6.3.1. *Let $P = \prod_{i=1}^{n} p_i$. Then,* $\mathrm{CS}_0 = P \cdot \mathbf{E}[Z_1]$.

Proof. Note that CS_0 is the expected surplus lost due to not allocating the item in Myerson's mechanism. This happens only when all realized values are below their corresponding reserve price. In this case, the value lost is the maximum valuation, since this value contributes to the welfare, and the revenue raised is zero. Therefore, we have:

$$\mathrm{CS}_0 = \left(\prod_{i=1}^{n} p_i\right) \cdot \mathbf{E}\left[\max_{i=1,2,\ldots,n} v_i \,\middle|\, \forall i, v_i < r_{\mathcal{D}_i}\right]$$

where the expectation is over $\vec{v} \sim \mathcal{D}$. This is equal to $P \cdot \mathbf{E}[Z_1]$. □

Vanilla Signaling Schemes. Before presenting our general signaling scheme, it is instructive to first consider simpler schemes to understand the challenges posed by this problem. One possible scheme is to pick a random buyer and apply the single-buyer BBM signaling scheme defined in Section 6.2 to it. Denote this buyer by i. Such a single-buyer signaling scheme will construct the set of BBM signals for buyer i by decomposing $\mathcal{D}_i$ and pretending the other buyers don't exist. Given the valuation $v_i \sim \mathcal{D}_i$ of this buyer, the scheme will send a signal $\mathrm{BBM}(v_i, \mathcal{D}_i)$ just as in single-buyer case, and reveal the identity of this buyer. There is no signal sent for the other buyers, so that the seller's information for $j \neq i$ is their prior $\mathcal{D}_j$.

The nice property of the BBM signaling scheme is that the virtual value of buyer i is always non-negative (Lemma 6.2.3). Therefore, in the event when all buyers $j \neq i$ have values $v_j < r_{\mathcal{D}_j}$ (so that their virtual values are negative), the seller allocates the item to i. Since v_i is independent of other buyers' values, the mechanism therefore behaves exactly as $\mathrm{BBM}(v_i, \mathcal{D}_i)$ from the perspective of buyer i. In other words, with probability $\prod_{j \neq i} p_j$, we generate the single-buyer CS from Lemma 6.2.3, which is at least:

$$\text{CS(BBM) from buyer } i \geq \Pr[v_i < r_{\mathcal{D}_i}]\, \mathbf{E}[v_i \mid v_i < r_{\mathcal{D}_i}] = p_i\, \mathbf{E}[Y_i].$$

Since buyer i is a randomly chosen buyer, the overall CS generated is:

$$\text{Overall CS} \geq \frac{1}{n}\sum_{i=1}^{n}\left(\prod_{j\neq i} p_j\right) p_i\, \mathbf{E}[Y_i] = \left(\prod_{j=1}^{n} p_j\right)\frac{\sum_{i=1}^{n}\mathbf{E}[Y_i]}{n} = P\frac{\sum_{\ell=1}^{n}\mathbf{E}[Z_\ell]}{n} \geq P\frac{\mathbf{E}[Z_1]}{n}$$

where $P = \prod_{j=1}^{n} p_j$. Therefore, comparing the expression above with that in Lemma 6.3.1, this scheme yields an n-approximation to CS_0. Further, we can construct identical regular distributions for which the expected value of the max is comparable to the expected value of the sum, that is, $\sum_{\ell=1}^{n}\mathbf{E}[Z_\ell] = O(\mathbf{E}[Z_1])$. Therefore, this analysis cannot be improved.

On the other hand, the above scheme does achieve CS $=$ Opt for two-valued i.i.d. distributions ($\mathcal{K} = 2$ and arbitrary n). To see this, assume the support is $a < b$, and let $q = \Pr[\mathcal{D}_i = a]$. If the reserve price is a, Myerson's auction is already efficient, that is, $\text{CS}(\mathcal{D}) = \text{Opt}$; else Myerson's auction has $\text{CS}(\mathcal{D}) = 0$. We now have $Z_i = a$ for all i and $p_i = q$, so that the above scheme has surplus $q^n a = q^n\, \mathbf{E}[Z_1] = \text{CS}_0 = \text{Opt}$. Therefore, in either case, we extract CS $=$ Opt. Interestingly, this also shows that our lower bounds in Theorems 6.1.1 and 6.1.2 do require non-i.i.d. distributions when each $\mathcal{K}_i = 2$ regardless of the number n of buyers.

Moving beyond the $\mathcal{K} = 2$ setting to general $\mathcal{K}$ and n, it is tempting to run the BBM signaling scheme directly on the buyer with highest value, hoping to extract surplus $P\,\mathbf{E}[Z_1]$. However, this requires revealing the identity of the highest buyer to the seller, since the signaling scheme itself is public knowledge. But if the seller knows the identity of the highest buyer, she can always increase the reserve price to be the second highest bid. In other words, the posterior of the highest buyer is *truncated* at Z_2. This case needs a more careful construction of the signal and analysis, since the event of a buyer being the largest and hence BBM being applied to it is now correlated with the surplus this buyer generates in BBM. We perform this construction and analysis in Lemma 6.3.3. Intuitively, signaling using BBM on the largest buyer will only yield CS $\geq P\,\mathbf{E}[Z_1 - Z_2]$, which is again an $\Omega(n)$-approximation to $P\,\mathbf{E}[Z_1]$.

Our signaling scheme in the next section chooses a middle ground between these extremes – we will choose a rank $t \in \{1, 2, \ldots, n\}$ carefully, and choose a buyer whose value lies in the top t ranks at random. We will then perform the single-buyer BBM scheme on this buyer as we describe below. Surprisingly, this improves the naïve n-approximation to an $O(1)$-approximation!

6.3.2 Ranking-Based Multi-Buyer Signaling Scheme

We now introduce the family of signaling schemes $\textsc{Rank}_t$. We will derive a lower bound for the CS obtained by these schemes in our key lemma, Lemma 6.3.3. We present this scheme assuming $\mathcal{D}_i$ can be non-identical for generality.

Recall the definitions of $p_i, r_{\mathcal{D}_i}, Y_i, Z_\ell$ from above. In order to define the signaling scheme, we need an additional definition. For agent i with $V_i \sim \mathcal{D}_i$, we use $\mathcal{D}_{i|>a}$ to denote the conditional distribution of V_i given $V_i > a$, and $\mathcal{D}_{i|<a}$ to denote the conditional distribution of V_i given $V_i < a$. Moreover, we use $\mathcal{D}_{i|>a} - b$ to denote the distribution of $V_i - b$ given $V_i > a$; we refer to it as the distribution of V_i truncated at a and reduced by b.

The following result relates the reserve price of the truncated and the original distributions.

Lemma 6.3.2. *Let $\mathcal{D}'_i = \mathcal{D}_{i|>v^\circ} - v^\circ$ for any v°. Then we have $r_{\mathcal{D}'_i} \geq r_{\mathcal{D}_i} - v^\circ$, and moreover, for any $v > v^\circ$, we have $\varphi_{\mathcal{D}'_i}(v - v^\circ) = \varphi_{\mathcal{D}_i}(v) - v^\circ$.*

Proof. We first prove the result about reserve prices. Let $r = r_{\mathcal{D}_i}$ and $r' = r_{\mathcal{D}'_i}$. If $r \leq v^\circ$ the inequality is trivial. Otherwise, suppose for contradiction that for some $r' < r - v^\circ$, $r' \cdot S_{\mathcal{D}'_i}(r') \geq (r - v^\circ) \cdot S_{\mathcal{D}'_i}(r - v^\circ)$. Then we would have: $(r' + v^\circ) \cdot S_{\mathcal{D}'_i}(r') \geq r \cdot S_{\mathcal{D}'_i}(r - v^\circ)$, since $S_{\mathcal{D}'_i}(r') \geq S_{\mathcal{D}'_i}(r - v^\circ)$. Thus, $(r' + v^\circ) \cdot S_{\mathcal{D}_i}(r' + v^\circ) \geq r \cdot S_{\mathcal{D}_i}(r)$, a contradiction to the assumption that r is the smallest optimal reserve price of $\mathcal{D}_i$.

To see the second part, if we condition the distribution on $v > v^\circ$, this does not change the virtual value function for values $v_j > v^\circ$, since both the numerator and denominator in Eq. (6.1) scale by the same amount. If we now subtract v° from the support of the distribution, it reduces the virtual value by the same amount. This completes the proof. $\square$

Note that $\mathcal{D}'_i = \mathcal{D}_{i|>v^\circ} - v^\circ$ represents buyer i's *excess value* compared to v°. Lemma 6.3.2 shows that for any threshold v° and any buyer i, given the side-information that $V_i > v^\circ$, her new reserve price is greater than her original reserve price.

The $\textsc{Rank}_t$ *signaling scheme.* We now present the family of signaling schemes $\textsc{Rank}_t$ parameterized by the rank $t \in \{1, \ldots, n\}$. For any realized joint valuation profile $\vec{v} = (v_1, v_2, \ldots, v_n)$, the signal sent by $\textsc{Rank}_t$ consists of two parts. In the first part, $\textsc{Rank}_t$ observes $\vec{v}$ and outputs (v°, T), where v° is the value of $(t+1)^{\text{st}}$ largest realized value (or 0 when $t = n$), and T is the subset of buyers with realized value strictly greater than v°. For the second part of the signal, $\textsc{Rank}_t$ chooses a buyer j uniformly at random from T, and computes her excess distribution $\mathcal{D}_{j|>v^\circ} - v^\circ$. It then reveals both the identity of j, as well as the signal $\text{BBM}(v_j - v^\circ, \mathcal{D}_{j|>v^\circ} - v^\circ)$ generated by the single-buyer BBM scheme on a buyer with value distribution $\mathcal{D}_{j|>v^\circ} - v^\circ$. The scheme is formalized in Algorithm 5.

Optimal mechanism under $\textsc{Rank}_t$. Conditioned on receiving the signal generated by $\textsc{Rank}_t$, the seller is guaranteed a revenue of v° from the $(t+1)^{\text{st}}$ largest buyer, and knows that only buyers in T can pay more than v°. The seller can now charge at least v° to any buyer in T, and can further run an auction over the excess value of buyers in T, where for buyer $i \in T$, her excess value has distribution $\mathcal{D}'_i = \mathcal{D}_{i|>v^\circ} - v^\circ$. Note that for any buyer $i \in T$ except the randomly chosen buyer j, a value drawn from $\mathcal{D}'_i$ represents how much more than v° she is willing to pay. Moreover, distributions $\mathcal{D}'_i$ are independent, and also, since the identity of j is chosen uniformly at random, the BBM scheme modifies the distribution of buyer j in a fashion that is independent of $\mathcal{D}'_i$.

By Lemma 6.2.3, we know that the BBM scheme ensures the virtual value of buyer j is always non-negative. From the characterization of the optimal auction [Mye81, Elk07], since the item is always allocated to the highest virtual value buyer as long as this value is non-negative, the item will always be allocated to buyer j if all other buyers $i \in T, i \neq j$ have excess values $v_i - v^\circ < r_{\mathcal{D}'_i}$ (and hence, negative virtual values).

Algorithm 5 $\textsc{Rank}_t(\vec{v}, \mathcal{D})$

```
v° ← (t+1)^st largest value in v⃗
T ← {i : v_i > v°}
if T ≠ φ then
    j ← Buyer chosen uniformly at random from T
    s ← BBM(v_j − v°, D_{j|>v°} − v°)
    return v°, T, j, and s
else
    return v°, and T = φ
end if
```

Consumer surplus under $\textsc{Rank}_t$. We require the following key lemma, that gives a lower bound for the consumer surplus generated under $\textsc{Rank}_t$. This lemma forms the crux of our subsequent analysis, helping us quantify how the BBM signal recovers much of CS lost by Myerson's auction. The difficulty in proving it arises because the random choice of buyer j in Algorithm 5 is correlated with its rank, which in turn is correlated with its winning the auction and the surplus it generates in BBM. We get around this correlation by carefully *coupling* the surplus generated when buyer values are above the reserve with the order statistics of buyer valuations below the reserve.

Lemma 6.3.3. *For $1 \leq t \leq n$, the consumer surplus of* Rank_t *satisfies:*

$$\text{CS}(\text{Rank}_t) \geq \left(\prod_{i=1}^{n} p_i\right) \cdot \left(\left(\frac{1}{t} \cdot \sum_{\ell=1}^{t} \mathbf{E}[Z_\ell]\right) - \mathbf{E}[Z_{t+1}]\right).$$

Proof. For convenience, denote $r_i = r_{\mathcal{D}_i}$. Fix a buyer b, and any valuation profile $\vec{v}_{-b} = \{v_i, i \neq b\}$ such that $v_i < r_i \, \forall \, i \neq b$. Define v_b^t as the t^{th} largest value in $\{v_i, i \neq b\}$. Now consider the event

$$Q(\vec{v}_{-b}, b, t) = \{V_b > v_b^t \text{ AND } b \text{ selected for BBM signaling}\}.$$

Conditioned on $Q(\vec{v}_{-b}, b, t)$, we have that the Rank_t scheme (Algorithm 5) with parameter t sets threshold value as $v^\circ = v_b^t$. By Lemma 6.3.2, we have that for every $i \in T, i \neq b$, their value v_i is smaller than their new reserve price $v^\circ + r_{\mathcal{D}'_i}$, since $v_i < r_{\mathcal{D}_i}$, and modifying $\mathcal{D}_i$ to $\mathcal{D}'_i = \mathcal{D}_{i|>v^\circ} - v^\circ$ does not decrease the reserve price. Therefore, conditioned on $Q(\vec{v}_{-b}, b, t)$, the auction behaves like the single item mechanism $\text{BBM}(v'_b, \mathcal{D}'_b)$, where $v'_b = v_b - v_b^t$, $\mathcal{D}'_b = \mathcal{D}_{b|>v_b^t} - v_b^t$. Let r'_b denote the reserve price of $\mathcal{D}'_b$; again using Lemma 6.3.2 we have $r'_b \geq r_b - v_b^t$. Now, using Lemma 6.2.3, we get that the expected consumer surplus generated by Rank_t under $Q(\vec{v}_{-b}, b, t)$ is at least:

$$\mathbf{E}[\text{CS}(\text{Rank}_t) \mid Q(\vec{v}_{-b}, b, t)] \geq \sum_{v'_b < r'_b} v'_b \Pr[\mathcal{D}'_b = v'_b] \geq \sum_{v_b^t < v_b < r_b} (v_b - v_b^t) \frac{f_{\mathcal{D}_b}(v_b)}{S_{\mathcal{D}_b}(v_b^t)}.$$

Note also that $\Pr[Q(\vec{v}_{-b}, b, t)] = \frac{1}{t} \cdot \left(\prod_{i \neq b} f_{\mathcal{D}_i}(v_i) \mathbb{1}_{\{v_i < r_{\mathcal{D}_i}\}}\right) \cdot S_{\mathcal{D}_b}(v_b^t)$. Thus for any buyer b, and any valuation profile $\vec{v}_{-b}$ with $v_i < r_i$ for all $i \neq b$, we have

$$\begin{aligned}
\mathbf{E}[\mathrm{CS}(\textsc{Rank}_t) \cdot \mathbb{1}_{Q(\vec{v}_{-b},b,t)}] &= \Pr[Q(\vec{v}_{-b}, b, t)] \cdot \mathbf{E}[\mathrm{CS}(\textsc{Rank}_t) \mid Q(\vec{v}_{-b}, b, t)] \\
&\geq \frac{1}{t} \left(\prod_{i \neq b} f_{\mathcal{D}_i}(v_i)\right) \sum_{v_b^t < v_b < r_b} (v_b - v_b^t) f_{\mathcal{D}_b}(v_b) \\
&= \frac{1}{t} \sum_{v_b < r_b} \left(\prod_{i} f_{\mathcal{D}_i}(v_i)\right) \max\{(v_b - v_b^t), 0\}.
\end{aligned}$$

For any $\vec{v}$ let $v^{(t)}$ denote the t^{th} largest value, and $I_t(\vec{v})$ to be the indices corresponding to the top t values in $\vec{v}$. Summing up over all b, and all $\vec{v}_{-b}$ such that $v_i < r_i \, \forall \, i \neq b$, we have

$$\begin{aligned}
\sum_b \sum_{\vec{v}_{-b}} \mathbf{E}[\mathrm{CS}(\textsc{Rank}_t) \cdot \mathbb{1}_{Q(\vec{v}_{-b},b,t)}] &\geq \sum_{\vec{v} | v_i < r_i} \frac{1}{t} \cdot \left(\prod_{i} f_{\mathcal{D}_i}(v_i)\right) \cdot \left(\sum_{b} \max\{(v_b - v_b^t), 0\}\right) \\
&= \sum_{\vec{v} | v_i < r_i} \left(\prod_{i} f_{\mathcal{D}_i}(v_i)\right) \left(\sum_{i \in I_t(\vec{v})} \frac{1}{t} \left(v_i - v^{(t+1)}\right)\right) \\
&= \sum_{\vec{v} | v_i < r_i} \left(\prod_{i} f_{\mathcal{D}_i}(v_i)\right) \left(\left(\sum_{i \in I_t(\vec{v})} \frac{v_i}{t}\right) - v^{(t+1)}\right).
\end{aligned} \tag{6.2}$$

Let $\mathcal{D}''_i = \mathcal{D}_{i|<r_{\mathcal{D}_i}}$ be the distribution of buyer i's value conditioned on $V_i < r_{\mathcal{D}_i}$. Recall we define $p_i = \Pr_{v_i \sim \mathcal{D}_i}[v_i < r_{\mathcal{D}_i}]$; thus $f_{\mathcal{D}''_i}(v) = f_{\mathcal{D}_i}(v)/p_i$. Suppose we independently sample $Y_i \sim \mathcal{D}_{i|<r_{\mathcal{D}_i}}$ for each i, and define Z_ℓ as the ℓ^{th} largest value

in $\{Y_i\}$. Then Eq. (6.2) can be written as

$$\sum_b \sum_{\vec{v}_{-b} | v_i < r_i \, \forall \, i \neq b} \mathbf{E}[\mathrm{CS}(\textsc{Rank}_t) \cdot \mathbb{1}_{Q(\vec{v}_{-b}, b, t)}]$$

$$\geq \sum_{\vec{v} | v_i < r_i} \left(\prod_i p_i \right) f_{\mathcal{D}''}(\vec{v}) \left(\left(\sum_{i \in I_t(\vec{v})} \frac{v_i}{t} \right) - v_{(t+1)} \right)$$

$$= \left(\prod_i p_i \right) \left(\left(\sum_{\ell=1}^{t} \frac{\mathbf{E}[Z_\ell]}{t} \right) - \mathbf{E}[Z_{t+1}] \right).$$

Finally, noting that the $Q(\vec{v}_{-b}, b, t)$ events are all non-overlapping, we can write

$$\mathrm{CS}(\textsc{Rank}_t) \geq \sum_b \sum_{\vec{v}_{-b} | v_i < r_i \, \forall \, i \neq b} \mathbf{E}[\mathrm{CS}(\textsc{Rank}_t) \cdot \mathbb{1}_{Q(\vec{v}_{-b}, b, t)}],$$

thereby completing the proof. $\square$

6.3.3 Approximating Non-allocation Surplus CS_0 *and Proof of Theorem 6.1.3*

Given the above signaling scheme, approximating CS_0 (and hence showing Theorem 6.1.3) is now simple: We choose the parameter $t \in \{1, 2, \ldots, n\}$ that maximizes $\mathrm{CS}(\textsc{Rank}_t)$ and run $\textsc{Rank}_t$. We denote this scheme as $\{\mathcal{S}_0(\vec{v}, \mathcal{D})\}_{\vec{v}}$, and present it in Algorithm 6.

Algorithm 6 $\mathcal{S}_0(\vec{v}, \mathcal{D})$

1: Choose $t = \mathrm{argmax}_{t'=1}^{n} \mathrm{CS}(\textsc{Rank}_{t'})$.
2: **return** $\textsc{Rank}_t(\vec{v}, \mathcal{D})$

The following theorem shows that this scheme approximates CS_0 when the $\mathcal{D}_i$ are a common regular distribution Θ. Since $\textsc{Opt} = \mathrm{CS}(\mathcal{D}) + \mathrm{CS}_0$ when $\mathcal{D}_i$ are regular and identical, Theorem 6.3.4 shows the better of no signaling and Algorithm 6 is an $O(1)$-approximation to $\textsc{Opt}$, completing the proof of Theorem 6.1.3.

Theorem 6.3.4. *The consumer surplus of the signaling scheme* $\mathcal{S}_0(\vec{v}, \mathcal{D})$ *is an* $O(1)$*-approximation to the non-allocation surplus,* CS_0*.*

Let $\textsc{Alg}$ denote the CS of $\mathcal{S}_0(\vec{v}, \mathcal{D})$. The above theorem will follow from the following, since $\textsc{Alg}$ is at least the LHS by Lemma 6.3.3 and CS_0 is equal to $P \cdot \mathbf{E}[Z_1]$ by Lemma 6.3.1.

Theorem 6.3.5. *When each $\mathcal{D}_i$ is i.i.d. and regular with common distribution Θ, then we have:*

$$P \cdot \max_t \left(\left(\frac{1}{t} \sum_{i=1}^{t} \mathbf{E}[Z_i] \right) - \mathbf{E}[Z_{t+1}] \right) \geq \frac{1}{1900} \cdot P \cdot \mathbf{E}[Z_1].$$

The rest of this section is devoted to proving Theorem 6.3.5. Note that the above theorem is not true if the distribution Y (the conditional distribution of Θ below its reserve) on which the order statistics Z_i are defined, is a generic regular distribution such as an Exponential distribution. Consider the following example:

Example 6.3.6. Suppose Y is `Exponential`(1). Then $\mathbf{E}[Z_i] = \sum_{j=i}^{n} \frac{1}{j} = H_n - H_{i-1}$, where $H_i = \sum_{j=1}^{i} \frac{1}{j}$. Therefore $\text{CS}_0 = \mathbf{E}[Z_i] = H_n$. It is easy to check that $\left(\frac{1}{t} \sum_{i=1}^{t} \mathbf{E}[Z_i]\right) - \mathbf{E}[Z_{t+1}] = 1 + H_n - H_t - (H_n - H_t) = 1$. Therefore, $\text{CS}_0 = \Omega(\log n) \cdot \textsc{Alg}$.

Note however that Y is of a more specific form: It is the conditional distribution of a regular distribution Θ below its reserve r. This means in particular that Y *cannot* be an Exponential distribution (or its discrete counterpart, the Geometric distribution). We will crucially use the property that the revenue of Θ when the price is set below the reserve r is a concave function of the quantile of the price, and further, this function is non-decreasing. Formally, we use the following.

Lemma 6.3.7. *For any regular distribution Θ with $S(v) = \Pr[\Theta \geq v]$ and reserve r, and for any $u \leq v < r$ in the support of Θ, we have: (1) $u \cdot S(u) \geq \frac{1-S(u)}{1-S(v)} \cdot v \cdot S(v)$, and (2) $u \cdot S(u) \leq v \cdot S(v)$.*

Proof. Let $f_i = \Pr[\mathcal{D} = i]$. Regularity means the quantity $\varphi(x_i) = x_i - (x_{i+1} - x_i)\frac{S(x_{i+1})}{f(x_i)}$ is non-decreasing in i. Let $R_i = x_i S(x_i)$ be the revenue when the price is x_i. Note that

$$\varphi(x_i) = \frac{R_i - R_{i+1}}{S(x_i) - S(x_{i+1})}.$$

Consider the revenue as a function of $S(x_i)$. Regularity means this function is concave on $S(x_i)$ where x_i is within the support of $\mathcal{D}$. The lemma then follows since the curve is non-increasing in $S(x_i)$ for $x_i \leq r$. $\square$

The next idea in the analysis is to decompose the distribution Y into the "core" and "tail", so that values v with $\Pr[Y \geq v] = O(1/n)$ lie in the tail, and the rest lie in the core. Roughly speaking, we will show that the expected value of Y in either the core or the tail is upper bounded by $\mathrm{CS}(\textsc{Rank}_t)$ for suitable choices of t. Such choices are non-trivial and form the crux of our analysis.

Core-Tail Decomposition

To fix notation, recall that every $\mathcal{D}_i$ is the same regular distribution Θ. The values $x_1 < x_2 < \cdots < x_{\mathcal{K}}$ form the support of Θ and its reserve price is r. Recall $S(x) = \Pr[\Theta \geq x]$ and Y is the conditional distribution of Θ strictly below r.

First, if $n \leq 1000$, $\textsc{Alg} \geq \frac{1}{1000}\mathbf{E}[Z_1] = \frac{1}{1000}\mathrm{CS}_0$. Therefore, we assume $n > 1000$ in the following analysis. We further assume $r > x_1$ since otherwise $\mathrm{CS}_0 = 0$ and Myerson's auction itself raises optimal CS. Suppose $r = x_i$ with $i > 1$, then $\{x_1, \ldots, x_{i-1}\}$ is the support of Y.

Let u^+ be the smallest value in the support of Y with $\Pr[Y \geq u^+] \leq \frac{4}{n}$ and let u^- be the largest value in the support of Y with $\Pr[Y \geq u^-] > \frac{4}{n}$. Both of these values exist since $\Pr[Y \geq x_1] = 1$ (we assumed $r > x_1$) and $\Pr[Y \geq r] = 0$. Further, u^- and u^+ are consecutive values in the support of Y.

We divide CS_0 into $\textsc{Core}$ and $\textsc{Tail}$, based on whether there is a realized value at least u^+.

Definition 6.3.8. $\textsc{Core} := P \cdot u^-$, and $\textsc{Tail} := P \cdot n \cdot \mathbf{E}[Y \cdot \mathbb{1}(Y \geq u^+)]$.

Lemma 6.3.9. *$\mathrm{CS}_0 \leq \textsc{Core} + \textsc{Tail}$.*

Proof. This comes immediately from

$$\begin{aligned}\mathrm{CS}_0 &= P \cdot (\mathbf{E}[Z_1 \cdot \mathbb{1}(Z_1 \leq u^-)] + \mathbf{E}[Z_1 \cdot \mathbb{1}(Z_1 \geq u^+)]) \\ &\leq P \cdot (u^- \cdot \Pr[Z_1 \leq u^-] + n \cdot \mathbf{E}[Y \cdot \mathbb{1}(Y \geq u^+)]) \\ &\leq \textsc{Core} + \textsc{Tail}.\end{aligned}$$

where the first inequality upper bounds the max of n i.i.d. samples from $\mathbb{1}(Y \geq u^+)$ by their sum. □

Upper Bound on $\textsc{Core}$ *and* $\textsc{Tail}$

We will separately bound $\textsc{Core}$ and $\textsc{Tail} - O(1) \cdot \textsc{Core}$ in terms of $\textsc{Alg}$. We first bound $\textsc{Tail}$ using the following lemma.

Lemma 6.3.10. *$\textsc{Alg} \geq \frac{1}{250} \cdot (\textsc{Tail} - 4 \cdot \textsc{Core})$.*

Proof. If $\Pr[Y \geq u^+] = 0$, then the inequality trivially holds since $\textsc{Tail} = 0$. Otherwise,

$$\begin{aligned}\textsc{Alg} \geq \mathrm{CS}(\textsc{Rank}_1) &= P \cdot \mathbf{E}[Z_1 - Z_2] \\ &\geq P \cdot \Pr[Z_1 \geq u^+ \wedge Z_2 \leq u^-] \cdot \mathbf{E}[Z_1 - Z_2 \mid Z_1 \geq u^+ \wedge Z_2 \leq u^-] \\ &\geq P \cdot \Pr[Z_1 \geq u^+ \wedge Z_2 \leq u^-] \cdot (\mathbf{E}[Z_1 \mid Z_1 \geq u^+ \wedge Z_2 \leq u^-] - u^-).\end{aligned}$$

Now, to bound the term $\Pr[Z_1 \geq u^+ \wedge Z_2 \leq u^-]$, we have the following. Note here

u^- and u^+ are consecutive values in the support of Y.

$$
\begin{aligned}
\Pr[Z_1 \geq u^+ \wedge Z_2 \leq u^-] &= \Pr[Z_1 \geq u^+] \cdot \frac{\Pr[Z_1 \geq u^+ \wedge Z_2 \leq u^-]}{\Pr[Z_1 \geq u^+]} \\
&= \Pr[Z_1 \geq u^+] \cdot \frac{n \cdot (1 - \Pr[Y \leq u^-]) \cdot \Pr[Y \leq u^-]^{n-1}}{1 - \Pr[Y \leq u^-]^n} \\
&= \Pr[Z_1 \geq u^+] \cdot \frac{n}{\sum_{i=0}^{n-1} \Pr[Y \leq u^-]^i} \cdot \Pr[Y \leq u^-]^{n-1} \\
&\geq \Pr[Z_1 \geq u^+] \cdot \Pr[Y \leq u^-]^{n-1} \\
&\geq \Pr[Y \leq u^-]^n = \Pr[Z_1 \geq u^+] \cdot \Pr[Z_1 \leq u^-].
\end{aligned}
$$

We now bound $\Pr[Z_1 \geq u^+]$ and $\Pr[Z_1 \leq u^-]$ separately. To bound $\Pr[Z_1 \geq u^+]$, note that

$$
\Pr[Z_1 \geq u^+] = 1 - (1 - \Pr[Y \geq u^+])^n \geq 1 - \exp(-n \cdot \Pr[Y \geq u^+]) \geq \frac{1 - e^{-4}}{4} \cdot n \cdot \Pr[Y \geq u^+]
$$

where the final inequality follows since the function $\psi(x) = \frac{1 - e^{-4x}}{(1-e^{-4})x} \geq 1$ when $x = \frac{1}{4} \cdot n \cdot \Pr[Y \geq u^+] \leq 1$. To bound $\Pr[Z_1 \leq u^-]$, note that:

$$
\Pr[Z_1 \leq u^-] = (1 - \Pr[Y \geq u^+])^n \geq (1 - 4/n)^n \geq 0.018
$$

where we have used $n \geq 1000$ in the last inequality. Therefore, we have

$$
\Pr[Z_1 \geq u^+ \wedge Z_2 \leq u^-] \geq \Pr[Z_1 \geq u^+] \cdot \Pr[Z_1 \leq u^-] \geq 0.004 \cdot n \cdot \Pr[Y \geq u^+].
$$

To bound $\textsc{Alg}$, we also need to simplify the term $\mathbf{E}[Z_1 \mid Z_1 \geq u^+ \wedge Z_2 \leq u^-]$. For this, let $Y_1, \ldots, Y_n$ be independent draws from Y. We have:

$$
\begin{aligned}
\mathbf{E}[Z_1 \mid Z_1 \geq u^+ \wedge Z_2 \leq u^-] &= \sum_{i=1}^{n} \frac{1}{n} \cdot \mathbf{E}[Y_i \mid Y_i \geq u^+ \wedge Y_j \leq u^-, \ \forall j \neq i] \\
&= \mathbf{E}[Y_1 \mid Y_1 \geq u^+ \wedge Y_j \leq u^-, \ \forall j \neq 1] \\
&= \mathbf{E}[Y_1 \mid Y_1 \geq u^+] = \mathbf{E}[Y \mid Y \geq u^+].
\end{aligned}
$$

Here, the first equality follows since any Y_i is equally likely to be the maximum value, and the third equality follows by the independence of Y_i's. Putting all this together, we bound ALG as:

$$\begin{aligned}
\textsc{Alg} &\geq P \cdot \Pr[Z_1 \geq u^+ \wedge Z_2 \leq u^-] \cdot (\mathbf{E}[Z_1 \mid Z_1 \geq u^+ \wedge Z_2 \leq u^-] - u^-) \\
&\geq 0.004P \cdot n \cdot \Pr[Y \geq u^+] \cdot (\mathbf{E}[Y \mid Y \geq u^+] - u^-) \\
&= 0.004P \cdot (n \cdot \mathbf{E}[Y \cdot \mathbb{1}(Y \geq u^+)] - n \cdot \Pr[Y \geq u^+] \cdot u^-) \\
&\geq 0.004 \cdot (\textsc{Tail} - 4 \cdot \textsc{Core}).
\end{aligned}$$

where the last inequality uses $n \cdot \Pr[Y \geq u^+] \leq 4$. $\square$

Next, we bound CORE, crucially using the regularity of Θ (via Lemma 6.3.7).

Lemma 6.3.11. *ALG $\geq \frac{1}{330} \cdot$ CORE.*

Proof. We divide the proof into four cases based on the value of $S(u^-)$.

Case (1): $S(u^-) \geq 0.1$ and there is a value w with $S(w) \in [0.75 \cdot S(u^-) + 0.25, 0.25 \cdot S(u^-) + 0.75]$. We have $w < u^- < r$. Note that $S(u^-) \geq 0.1$ and $S(w) \leq 0.25 \cdot S(u^-) + 0.75$ together imply $S(u^-) \geq \frac{0.1}{0.775} \cdot S(w)$. Also note that $S(w) \geq 0.75 \cdot S(u^-) + 0.25$ implies $1 - S(w) \geq \frac{1}{4}(1 - S(u^-))$. Using Lemma 6.3.7, we now have

$$w \cdot S(w) \geq \frac{1 - S(w)}{1 - S(u^-)} \cdot u^- \cdot S(u^-) \geq \frac{1}{4} \cdot u^- \cdot \frac{0.1}{0.775} \cdot S(w).$$

so that $w \geq \frac{1}{31} \cdot u^-$. To bound the performance of ALG, we have

$$\begin{aligned}
\textsc{Alg} \geq \mathrm{CS}(\textsc{Rank}_n) &= P \cdot \frac{1}{n} \sum_{i=1}^{n} \mathbf{E}[Z_i] \geq P \cdot \frac{\lfloor n/10 \rfloor}{n} \cdot \mathbf{E}[Z_{\lfloor n/10 \rfloor}] \\
&\geq P \cdot \frac{\lfloor n/10 \rfloor}{n} \cdot w \cdot \Pr[Z_{\lfloor n/10 \rfloor} \geq w] \\
&\geq P \cdot (1 - \Pr[Z_{\lfloor n/10 \rfloor} < w]) \cdot \frac{1}{320} \cdot u^- \\
&= \frac{1}{320} \cdot (1 - \Pr[Z_{\lfloor n/10 \rfloor} < w]) \cdot \textsc{Core}.
\end{aligned}$$

where we have used $w \geq \frac{1}{31} \cdot u^-$ and $n > 1000$.

Now notice that $\Pr[Y \geq w] \geq \Pr[\Theta \geq w] = S(w) \geq 0.25$. Consider the sum of n i.i.d. samples from $\texttt{Bernoulli}(1, 0.25)$. The quantity $\Pr[Z_{\lfloor n/10 \rfloor} < w]$ is upper bounded by the probability that the above sum is at most $\lfloor n/10 \rfloor$. We apply Chernoff bounds [MU05] to this Bernoulli sum to obtain:

$$\Pr[Z_{\lfloor n/10 \rfloor} < w] \leq \exp\left(-\frac{(0.25n - 0.1n)^2}{2 \cdot (0.25n)}\right) \leq \exp\left(-0.045 \cdot n\right) < 0.001.$$

where we have used $n > 1000$.

We conclude that $\textsc{Alg} \geq \frac{1}{330} \cdot \textsc{Core}$ in this case.

Case (2): $S(u^-) < 0.1$ and there is a value w with $S(w) \in [2 \cdot S(u^-), 4 \cdot S(u^-)]$. As before, we have $w < u^- < r$. Further, note that $S(u^-) < 0.1$ and $S(w) \leq 4 \cdot S(u^-)$ together imply $1 - S(w) \geq \frac{0.6}{0.9}(1 - S(u^-))$ and $S(u^-) \geq \frac{1}{4} S(w)$. Using Lemma 6.3.7, we have

$$w \cdot S(w) \geq \frac{1 - S(w)}{1 - S(u^-)} \cdot u^- \cdot S(u^-) \geq \frac{0.6}{0.9} \cdot u^- \cdot \frac{1}{4} \cdot S(w).$$

so that $w \geq \frac{1}{6} \cdot u^-$.

Let b be the largest value within the support of Y with $\Pr[Y \geq b] \geq 2 \cdot \Pr[Y \geq w]$. This implies $\Pr[Y > b] < 2 \Pr[Y \geq w]$. Note that such a value b exists since $\Pr[Y \geq x_1] = 1$. Since $b < w < r$, by Lemma 6.3.7, we have $b \cdot S(b) \leq w \cdot S(w)$ and thus

$$b \leq w \cdot \frac{S(w)}{S(b)} = w \cdot \frac{(S(w) - S(r)) + S(r)}{(S(b) - S(r)) + S(r)}.$$

Further, using the fact that $\frac{S(w) - S(r)}{S(b) - S(r)} = \frac{\Pr[Y \geq w]}{\Pr[Y \geq b]} \leq \frac{1}{2}$, and $S(w) - S(r) \geq S(r)$ (this is since $S(w) \geq 2 \cdot S(u^-)$), we get $b \leq \frac{2}{3} \cdot w$ and thus

$$w - b \geq \frac{1}{3} \cdot w \geq \frac{1}{18} \cdot u^-.$$

Choose $t \geq 40$ so that $\Pr[Y \geq w] \in [0.2 \cdot t/n, 0.3 \cdot t/n]$. We now show that such a t exists. First note that these intervals overlap for consecutive t. Next note that $\Pr[Y \geq w] \geq 2 \cdot \Pr[Y \geq u^-] \geq \frac{8}{n}$ (by the choice of u^-), so that when $t = 40$, the value $0.2 \cdot t/n$ is a lower bound on $\Pr[Y \geq w]$. Further, since $\Pr[Y \geq w] \leq 4 \Pr[Y \geq u^-] \leq 0.4$, this means that when $t = n$, the value $0.3 \cdot t/n$ is an upper bound on $\Pr[Y \geq w]$. This shows such a t must exist.

Putting these together, we now have:

$$\begin{aligned}
\textsc{Alg} \geq \mathrm{CS}(\textsc{Rank}_t) = P \cdot \frac{1}{t}\sum_{i=1}^{t} \mathbf{E}[Z_i - Z_{t+1}] &\geq P \cdot \frac{\lfloor t/10 \rfloor}{t} \cdot \mathbf{E}[Z_{\lfloor t/10 \rfloor} - Z_{t+1}] \\
&\geq P \cdot \frac{\lfloor t/10 \rfloor}{t} \cdot \Pr[Z_{\lfloor t/10 \rfloor} \geq w \wedge Z_t \leq b] \cdot (w - b) \\
&\geq P \cdot (1 - \Pr[Z_{\lfloor t/10 \rfloor} < w] - \Pr[Z_t > b]) \cdot \frac{1}{190} \cdot u^- \\
&= \frac{1}{190} \cdot (1 - \Pr[Z_{\lfloor t/10 \rfloor} < w] - \Pr[Z_t > b]) \cdot \textsc{Core}.
\end{aligned}$$

where we have used $w - b \geq \frac{1}{18} \cdot u^-$ and $n > 1000$. Now notice that $\Pr[Y \geq w] \geq 0.2 \cdot t/n$ and $\Pr[Y > b] < 2S_Y(w) \leq 0.6 \cdot t/n$. We apply Chernoff bounds as in Case (1) to obtain:

$$\Pr[Z_{\lfloor t/10 \rfloor} < w] \leq \exp\left(-\frac{(0.2t - 0.1t)^2}{2 \cdot (0.2t)}\right) \leq \exp\left(-0.025 \cdot t\right) < 0.368.$$

and

$$\Pr[Z_t > b] \leq \exp\left(-\frac{(t - 0.6t)^2}{(1 + 1/0.6) \cdot (0.6t)}\right) \leq \exp\left(-0.1 \cdot t\right) < 0.019.$$

where we have used $t \geq 40$. Combining this with the lower bound on $\textsc{Alg}$, we derive:

$$\textsc{Alg} \geq \frac{1}{190} \cdot (1 - \Pr[Z_{\lfloor t/10 \rfloor} < w] - \Pr[Z_t > b]) \cdot \textsc{Core} \geq \frac{1}{310} \cdot \textsc{Core}.$$

Case (3): $S(u^-) \geq 0.1$ and there is no value w with $S(w) \in [0.75 \cdot S(u^-) + 0.25, 0.25 \cdot S(u^-) + 0.75]$. Let w^+ be the smallest value within the support of Θ with $S(w^+) <$

$0.75 \cdot S(u^-) + 0.25$, and let w^- be the largest value within the support of Θ with $S(w^-) > 0.25 \cdot S(u^-) + 0.75$. Using the same argument as in the previous cases, notice that $w^+ \geq \frac{1}{5} \cdot u^-$, since

$$w^+ \cdot S(w^+) \geq \frac{1 - S(w^+)}{1 - S(u^-)} \cdot u^- \cdot S(u^-) \geq \frac{3}{4} \cdot u^- \cdot \frac{0.1}{0.325} \cdot S(w^+).$$

Furthermore, w^- and w^+ are consecutive in the support of Y and both values are less than r.

If $\Pr[Y \geq w^+] > 0.25$, the proof is similar to Case (1) that

$$\begin{aligned}
\textsc{Alg} \geq \mathrm{CS}(\textsc{Rank}_n) &\geq P \cdot \Pr[Z_{\lfloor n/10 \rfloor} \geq w^+] \cdot \frac{\lfloor n/10 \rfloor}{n} \cdot w^+ \\
&\geq P \cdot (1 - \Pr[Z_{\lfloor n/10 \rfloor} < w^+]) \cdot \frac{1}{60} \cdot u^- \\
&\geq \frac{1}{60} \cdot (1 - \Pr[Z_{\lfloor n/10 \rfloor} < w^+]) \cdot \textsc{Core},
\end{aligned}$$

and by Chernoff bounds analogous to Case (1) we have,

$$\Pr[Z_{\lfloor n/10 \rfloor} < w^+] \leq \exp\left(-\frac{(0.25n - 0.1n)^2}{2 \cdot (0.25n)}\right) \leq \exp\left(-0.045 \cdot n\right) < 0.1.$$

Now assume $\Pr[Y \geq w^+] \leq 0.25$. Choose $t \geq 20$ so that $\Pr[Y \geq w^+] \in [0.2 \cdot t/n, 0.25 \cdot t/n]$. Using the same argument as for Case (2), this t exists since $\Pr[Y \geq w^+] \geq \Pr[Y \geq u^-] \geq \frac{4}{n}$ and $\Pr[Y \geq w^+] \leq 0.25$. Similar to Case (2), we have

$$\begin{aligned}
\textsc{Alg} &\geq P \cdot \Pr[Z_{\lfloor t/10 \rfloor} \geq w^+ \wedge Z_t \leq w^-] \cdot \frac{\lfloor t/10 \rfloor}{t} \cdot (w^+ - w^-) \\
&\geq P \cdot (1 - \Pr[Z_{\lfloor t/10 \rfloor} < w^+] - \Pr[Z_t > w^-]) \cdot \frac{\lfloor t/10 \rfloor}{t} \cdot (w^+ - w^-) \\
&= P \cdot (1 - \Pr[Z_{\lfloor t/10 \rfloor} < w^+] - \Pr[Z_t \geq w^+]) \cdot \frac{\lfloor t/10 \rfloor}{t} \cdot (w^+ - w^-).
\end{aligned}$$

where we have used that w^+ and w^- are consecutive values in the support of Y.

Now notice that $\Pr[Y \geq w^+] \geq 0.2 \cdot t/n$ and $\Pr[Y \geq w^+] \leq 0.25 \cdot t/n$ and as in Case (1), we apply Chernoff bounds as follows:

$$\Pr[Z_{\lfloor t/10 \rfloor} < w^+] \leq \exp\left(-\frac{(0.2t - 0.1t)^2}{2 \cdot (0.2t)}\right) \leq \exp\left(-0.025 \cdot t\right) < 0.607.$$

and

$$\Pr[Z_t \geq w^+] \leq \exp\left(-\frac{(t - 0.25t)^2}{(1 + 1/0.25) \cdot (0.25t)}\right) \leq \exp\left(-0.45 \cdot t\right) < 0.001.$$

Therefore, $\textsc{Alg} \geq \frac{1}{45} \cdot P \cdot (w^+ - w^-)$.

On the other hand, we have

$$\textsc{Alg} \geq \mathrm{CS}(\textsc{Rank}_n) \geq P \cdot \Pr[Z_{\lfloor n/2 \rfloor} \geq w^-] \cdot \frac{\lfloor n/2 \rfloor}{n} \cdot w^-.$$

Notice that $\Pr[Y \geq w^-] = \frac{S(w^-) - S(r)}{1 - S(r)} \geq \frac{S(w^-) - S(u^-)}{1 - S(u^-)} \geq 0.75$ and as in Case (1), we apply Chernoff bounds as follows:

$$\Pr[Z_{\lfloor n/2 \rfloor} < w^-] \leq \exp\left(-\frac{(0.75n - 0.5n)^2}{2 \cdot (0.75n)}\right) \leq \exp\left(-0.041 \cdot n\right) < 0.01.$$

Therefore, $\textsc{Alg} \geq \frac{1}{3} \cdot P \cdot w^-$ and further $\textsc{Alg} \geq \frac{1}{50} \cdot P \cdot w^+ \geq \frac{1}{250} \cdot \textsc{Core}$.

Case (4): $S(u^-) < 0.1$ and there is no value w with $S(w) \in [2 \cdot S(u^-), 4 \cdot S(u^-)]$. Let w^+ be the smallest value with $S(w^+) < 2 \cdot S(u^-)$, and let w^- be the largest value with $S(w^-) > 4 \cdot S(u^-)$. These values are consecutive in the support of Y.

We have $w^+ - w^- \geq \frac{1}{5} \cdot u^-$, since by Lemma 6.3.7,

$$w^+ \cdot S(w^+) \geq \frac{1 - S(w^+)}{1 - S(u^-)} \cdot u^- \cdot S(u^-) \geq \frac{0.8}{0.9} \cdot u^- \cdot \frac{1}{2} \cdot S(w^+),$$

and

$$w^+ \cdot S(w^+) \geq w^- \cdot S(w^-) \geq w^- \cdot 2 \cdot S(w^+).$$

where we have used $w^+ > w^-$. These imply $w^+ \geq \frac{4}{9} \cdot u^-$ and $w^- \leq \frac{w^+}{2}$, so that $w^+ - w^- \geq \frac{1}{5} \cdot u^-$.

Choose $t \geq 20$ so that $\Pr[Y \geq w^+] \in [0.2 \cdot t/n, 0.25 \cdot t/n]$. This t exists since $\Pr[Y \geq w^+] \geq \Pr[Y \geq u^-] \geq \frac{4}{n}$ and $\Pr[Y \geq w^+] \leq 0.2$. Similar to Case (2), this now gives

$$\begin{aligned}
\textsc{Alg} \geq \mathrm{CS}(\textsc{Rank}_t) &\geq P \cdot \Pr[Z_{\lfloor t/10 \rfloor} \geq w^+ \wedge Z_t \leq w^-] \cdot \frac{\lfloor t/10 \rfloor}{t} \cdot (w^+ - w^-) \\
&\geq P \cdot (1 - \Pr[Z_{\lfloor t/10 \rfloor} < w^+] - \Pr[Z_t > w^-]) \cdot \frac{1}{75} \cdot u^- \\
&\geq \frac{1}{75} \cdot (1 - \Pr[Z_{\lfloor t/10 \rfloor} < w^+] - \Pr[Z_t \geq w^+]) \cdot \textsc{Core}.
\end{aligned}$$

Now notice that $\Pr[Y \geq w^+] \geq 0.2 \cdot t/n$ and $\Pr[Y \geq w^+] < 0.25 \cdot t/n$. As in Case (1), we apply Chernoff bounds:

$$\Pr[Z_{\lfloor t/10 \rfloor} < w] \leq \exp\left(-\frac{(0.2t - 0.1t)^2}{2 \cdot (0.2t)}\right) \leq \exp(-0.025 \cdot t) < 0.607.$$

and

$$\Pr[Z_t > b] \leq \exp\left(-\frac{(t - 0.25t)^2}{(1 + 1/0.25) \cdot (0.25t)}\right) \leq \exp(-0.45 \cdot t) < 0.001.$$

Therefore, $\textsc{Alg} \geq \frac{1}{250} \cdot \textsc{Core}$. $\square$

Completing Proof of Theorem 6.3.5

Using Lemma 6.3.9, Lemma 6.3.10, and Lemma 6.3.11, we obtain:

$$\mathrm{CS}_0 \leq \textsc{Core} + \textsc{Tail} = 5 \cdot \textsc{Core} + (\textsc{Tail} - 4 \cdot \textsc{Core}) \leq 1650 \cdot \textsc{Alg} + 250 \cdot \textsc{Alg} = 1900 \cdot \textsc{Alg}.$$

This completes the proof of Theorem 6.3.5 and hence, Theorem 6.3.4 and Theorem 6.1.3.

7

Conclusions

Computational economics is a rapidly growing interdisciplinary subject that combines models and ideas in computer science and economics. In this dissertation, we demonstrated how to use techniques from computer science to approximate efficiency and fairness – primary objectives in the study of welfare economics. We looked into several developing theories in social choice and mechanism design, and provided foundations and breakthroughs in them.

There are countless economic models where exact efficiency and fairness are inaccessible, and guaranteed approximations to them are desirable. We believe the framework used in this dissertation can see its application in more settings in the coming years.

A

Omitted Proofs

A.1 The Running Time of Local Search

We now show that Algorithm 3 runs in polynomial time. This requires showing that the local search procedure in Section 3.4.2 runs in polynomial time; the rest of the steps in Algorithm 3 can easily be implemented efficiently. We will show this in two parts: The partial derivatives $\frac{\partial \phi_i(\mathbf{x})}{\partial x_j}$ can be approximately computed efficiently; and the number of iterations (finding the candidate pair (j, ℓ) and updating the allocations) performed by the local search procedure is polynomially bounded.

Estimating the Gradient of the Nash Welfare Objective

We first show how to estimate the derivative of $\phi(\mathbf{x})$ in the procedure in Section 3.4.2. First, the estimation procedure for U_i and its derivative is the same as that in [CCPV11]: Each time we compute $U_i(\mathbf{x})$, we pretend $\mathbf{x}$ is a product distribution over candidates, and sample it H times. Denote the samples as the committees $O_h \sim \mathbf{x}$ where $h \in \{1, 2, \ldots, H\}$. We calculate the quantity $\hat{U}_i(\mathbf{x}) = \frac{1}{H} \cdot \sum_{h=1}^{H} u_i(O_h)$ as the estimation of $U_i(\mathbf{x})$. The following lemma gives a bound on the additive error of estimation on $U_i(\mathbf{x})$ if H is sufficiently large.

Lemma A.1.1 ([CCPV11])**.** $\Pr\left[\left|\hat{U}_i(\mathbf{x}) - U_i(\mathbf{x})\right| > \Delta\right] < 2e^{-\frac{\Delta^2 \cdot H}{m^2}}$.

Since the derivative of U_i is given in Lemma 3.4.3, its error is bounded by estimating the multilinear function twice. This yields the following lemma:

Lemma A.1.2. *In the local search procedure in Section 3.4.2, suppose the total number of times the multilinear function U_i and its derivative is evaluated is N. Then, using poly$\left(N, m, \frac{1}{\Delta}\right)$ samples, the additive error in each estimate is bounded by Δ with probability $1 - \frac{1}{poly(m,N)}$.*

In the sequel, we will set $\Delta = \frac{\varepsilon^6}{64n \cdot m^5}$. We denote the estimated derivative of the multilinear function as $\frac{\hat{\partial} U_i(\mathbf{x})}{\hat{\partial} x_i}$, and the estimated derivative of $\phi_i(\mathbf{x})$ and $\phi(\mathbf{x})$ as $\frac{\hat{\partial}\phi_i(\mathbf{x})}{\hat{\partial} x_j} = \frac{1}{\hat{U}_i} \cdot \frac{\hat{\partial} U_i(\mathbf{x})}{\hat{\partial} x_j}$ and $\frac{\hat{\partial}\phi(\mathbf{x})}{\hat{\partial} x_j} = \sum_{i=1}^n \frac{\hat{\partial}\phi_i(\mathbf{x})}{\hat{\partial} x_j}$ respectively. We have the following lemma for bounding the error of estimating $\frac{\partial\phi(\mathbf{x})}{\partial x_j} \cdot \frac{1}{s_j}$. For this proof and subsequent ones, we crucially need that the Nash Welfare program in Section 3.4.2 lower bounds the allocations as $x_j \geq \underline{x}_j$.

Lemma A.1.3. *Let $\Delta = \frac{\varepsilon^6}{64n \cdot m^5}$, and suppose $x_j \geq \underline{x}_j$ for all candidates $j \in C_\ell$, where $\underline{x}_j$ is as defined in Section 3.4.2. If for any $j \in C_\ell$ and all $i \in V$, we have: $\left|\hat{U}_i(\mathbf{x}) - U_i(\mathbf{x})\right| < \Delta$ and $\left|\frac{\hat{\partial} U_i(\mathbf{x})}{\hat{\partial} x_j} - \frac{\partial U_i(\mathbf{x})}{\partial x_j}\right| < \Delta$, then it holds that $\frac{1}{s_j} \cdot \left|\frac{\hat{\partial}\phi(\mathbf{x})}{\hat{\partial} x_j} - \frac{\partial\phi(\mathbf{x})}{\partial x_j}\right| \leq \frac{\varepsilon}{8b}$.*

Proof. We have

$$\begin{aligned}
\left|\frac{\hat{\partial}\phi_i(\mathbf{x})}{\hat{\partial} x_j} - \frac{\mathsf{P}\phi_i(\mathbf{x})}{\mathsf{P}x_j}\right| &= \left|\frac{1}{\hat{U}_i(\mathbf{x})} \cdot \frac{\hat{\partial} U_i(\mathbf{x})}{\hat{\partial} x_j} - \frac{1}{U_i(\mathbf{x})} \cdot \frac{\mathsf{P}U_i(\mathbf{x})}{\mathsf{P}x_j}\right| \\
&\leq \frac{1}{U_i(\mathbf{x}) - \Delta} \cdot \left(\frac{\partial U_i(\mathbf{x})}{\partial x_j} + \Delta\right) - \frac{1}{U_i(\mathbf{x})} \cdot \frac{\mathsf{P}U_i(\mathbf{x})}{\mathsf{P}x_j} \\
&\leq \left|\frac{\Delta \cdot \left(U_i(\mathbf{x}) + \frac{\mathsf{P}U_i(\mathbf{x})}{\mathsf{P}x_j}\right)}{(U_i(\mathbf{x}) - \Delta)^2}\right|.
\end{aligned}$$

By Lemma 3.4.3, we have $\frac{\mathsf{P}U_i(\mathbf{x})}{\mathsf{P}x_j} \leq 1$. If $\Delta \leq \frac{\varepsilon^2}{2m^2}$, then by Eq. (3.6) we have $2\Delta \leq U_i(\mathbf{x})$, thus:

$$\left|\frac{\hat{\partial}\phi_i(\mathbf{x})}{\hat{\partial}x_j} - \frac{\mathsf{P}\phi_i(\mathbf{x})}{\mathsf{P}x_j}\right| = \left|\frac{\Delta \cdot \left(U_i(\mathbf{x}) + \frac{\mathsf{P}U_i(\mathbf{x})}{\mathsf{P}x_j}\right)}{(U_i(\mathbf{x}) - \Delta)^2}\right| \leq \left|\frac{\Delta \cdot (U_i(\mathbf{x}) + 1)}{U_i(\mathbf{x})^2/4}\right| \leq \left|\frac{\Delta \cdot 2}{\frac{\varepsilon^4}{m^4}/4}\right| \leq \left|\frac{8\Delta \cdot m^4}{\varepsilon^4}\right|.$$

Since $s_j \geq \frac{\varepsilon}{m} b$ for $j \in C_\ell$, we have

$$\frac{1}{s_j} \cdot \left|\frac{\partial\phi(\mathbf{x})}{\partial x_j} - \frac{\hat{\partial}\phi(\mathbf{x})}{\hat{\partial}x_j}\right| \leq \frac{1}{s_j}\sum_{i=1}^{n}\left|\frac{\hat{\partial}\phi_i(\mathbf{x})}{\hat{\partial}x_j} - \frac{\mathsf{P}\phi_i(\mathbf{x})}{\mathsf{P}x_j}\right| \leq \left|\frac{8\Delta \cdot m^4}{\varepsilon^4}\right| \cdot n \cdot \frac{m}{b \cdot \varepsilon} \leq \frac{8\Delta \cdot m^5 \cdot n}{\varepsilon^5 \cdot b} = \frac{\varepsilon}{8b},$$

finishing the proof. □

Number of Iterations in Local Search

The local search procedure in Section 3.4.2 iteratively finds a pair of candidates (j, ℓ) such that

$$\frac{\hat{\partial}\phi(\mathbf{x})}{\hat{\partial}x_j} \cdot \frac{1}{s_j} > \frac{\hat{\partial}\phi(\mathbf{x})}{\hat{\partial}x_\ell} \cdot \frac{1}{s_\ell} + \frac{3\varepsilon}{4b}. \tag{A.1}$$

Each time it finds such a pair, it increases x_j by $\delta \cdot \frac{1}{s_j}$ and decreases x_ℓ by $\delta \cdot \frac{1}{s_\ell}$. Note that at stopping, Eq. (A.1) does not hold, which by Lemma A.1.3 implies Eq. (3.5) also does not hold with high probability. Therefore, the termination condition in Section 3.4.2 is satisfied.

In order to bound the number of iterations of this procedure, we first show that the first and second partial derivatives of ϕ are bounded from above.

Lemma A.1.4. *If $\phi_i(\mathbf{x}) = \log U_i(\mathbf{x})$ and $x_j \geq \underline{x}_j$ for all candidates $j \in C_\ell$, where $\underline{x}_j$ is as defined in Section 3.4.2, then we have the following bounds for all $j, k \in C_\ell$ and $i \in W$:*

$$0 \leq \frac{\partial\phi_i(\mathbf{x})}{\partial x_j} = O\left(\frac{m^2}{\varepsilon^2}\right) \qquad \text{and} \qquad \left|\frac{\partial^2\phi_i(\mathbf{x})}{\partial x_j \partial x_k}\right| = O\left(\frac{m^4}{\varepsilon^4}\right).$$

Proof. Since$\phi_i(\mathbf{x}) = \log U_i(\mathbf{x})$, using Lemma 3.4.3, we have $\frac{\partial \phi_i(\mathbf{x})}{\partial x_j} \leq \frac{1}{U_i}$. Combining this with Eq. (3.6) gives $\frac{\mathsf{P}\phi_i(\mathbf{x})}{\mathsf{P}x_j} \leq \frac{5m^2}{\varepsilon^2}$.

We next bound the second order derivatives as follows. Here, $\mathbf{x_{-j,k}}$ is $\mathbf{x}$ with the j^{th} and k^{th} dimension removed, and $T \sim \mathbf{x_{-j,k}}$ means that T is chosen by including $\ell \in C \setminus \{j,k\}$ independently with probability x_ℓ. The first inequality below uses Eq. (3.6).

$$
\begin{aligned}
&\left|\frac{\partial^2 \phi_i(\mathbf{x})}{\partial x_j \partial x_k}\right| \\
=&\left|-\frac{1}{U_i(\mathbf{x})^2} \cdot \frac{\partial U_i(\mathbf{x})}{\partial x_j} + \frac{1}{U_i(\mathbf{x})} \cdot \frac{\partial^2 U_i(\mathbf{x})}{\partial x_j \partial x_k}\right| \\
\leq&\frac{25m^4}{\varepsilon^4} + \left|\frac{5m^2}{\varepsilon^2} \cdot \sum_{T \subseteq C \setminus \{j,k\}} \Pr_{T \sim \mathbf{x_{-j,k}}}[T] \cdot (u_i(T \cup \{j,k\}) - u_i(T \cup \{j\}) - u_i(T \cup \{k\}) + u_i(T))\right| \\
\leq&\frac{25m^4}{\varepsilon^4} + \frac{5m^2}{\varepsilon^2} \cdot \sum_{T \subseteq C \setminus \{j,k\}} \Pr_{T \sim \mathbf{x_{-j,k}}}[T] \cdot \Big(u_i(T \cup \{k\}) - u_i(T) - u_i(T \cup \{j,k\}) + u_i(T \cup \{j\})\Big) \\
\leq&\frac{25m^4}{\varepsilon^4} + \frac{5m^2}{\varepsilon^2} \cdot \sum_{T \subseteq C \setminus \{j,k\}} \Pr_{T \sim \mathbf{x_{-j,k}}}[T] \cdot 1 \leq \frac{26m^4}{\varepsilon^4}.
\end{aligned}
$$

This finishes the proof. □

The next lemma now follows from a standard application of first order Taylor approximation. We assume the local search procedure in Section 3.4.2 iteratively finds a pair of candidates (j, ℓ) such that Eq. (A.1) holds.

Lemma A.1.5. *Assuming all the estimates on $\frac{\mathsf{P}\phi(\mathbf{x})}{\mathsf{P}x_j}$ during execution of local search are within $\pm\frac{\varepsilon}{8b}$ of the true values, the total number of iterations is poly$\left(m, n, \frac{1}{\varepsilon}\right)$.*

Proof. Since all estimates of $\frac{\mathsf{P}\phi(\mathbf{x})}{\mathsf{P}x_j}$ during execution of local search are within $\pm\frac{\varepsilon}{8b}$ of the true values, if Eq. (A.1) holds for $j, k \in C_\ell$, then:

$$\frac{\mathsf{P}\phi(\mathbf{x})}{\mathsf{P}x_j} \cdot \frac{1}{s_j} > \frac{\hat{\partial}\phi(\mathbf{x})}{\hat{\partial}x_j} \cdot \frac{1}{s_j} - \frac{\varepsilon}{8b} \geq \frac{\hat{\partial}\phi(\mathbf{x})}{\hat{\partial}x_\ell} \cdot \frac{1}{s_\ell} + \frac{3\varepsilon}{4b} - \frac{\varepsilon}{8b} \geq \frac{\mathsf{P}\phi(\mathbf{x})}{\mathsf{P}x_\ell} \cdot \frac{1}{s_\ell} + \frac{\varepsilon}{2b}.$$

We next lower bound the increase in $\phi(\mathbf{x})$ when $\mathbf{x} = (\mathbf{x}_{-\mathbf{j},\mathbf{k}}, x_j, x_k)$ moves to $\mathbf{x}' = \left(\mathbf{x}_{-\mathbf{j},\mathbf{k}}, x_j + \frac{\delta}{s_j}, x_\ell - \frac{\delta}{s_\ell}\right)$. By Taylor approximation, we have:

$$\phi(\mathbf{x}') - \phi(\mathbf{x}) \geq (\mathbf{x}' - \mathbf{x}) \cdot \nabla\phi(\mathbf{x}) - \sum_{r=1}^{m} (x'_r - x_r)^2 \frac{M}{2}$$

where M is an upper bound on the absolute value of the second derivatives of ϕ. Since $M = O\left(\frac{m^4}{\varepsilon^4}\right)$ by Lemma A.1.4, this implies:

$$\phi(\mathbf{x}') - \phi(\mathbf{x}) \geq \delta \cdot \left(\frac{\mathsf{P}\phi(\mathbf{x})}{x_j} \cdot \frac{1}{s_j} - \frac{\mathsf{P}\phi(\mathbf{x})}{x_\ell} \cdot \frac{1}{s_\ell}\right) - \frac{26m^4}{\varepsilon^4} \cdot \frac{3\delta^2}{s_{\min}^2},$$

where $s_{\min} = \min_{\ell \in C_\ell} s_\ell \geq \frac{\varepsilon}{m} b$ and $\delta = \frac{\varepsilon^7 b}{312 m^6}$. Plugging these values in and simplifying, we have:

$$\phi(\mathbf{x}') - \phi(\mathbf{x}) = \Omega\left(\frac{\varepsilon\delta}{b}\right) = \text{poly}\left(\varepsilon, \frac{1}{n}, \frac{1}{m}\right).$$

Since $\phi(\mathbf{x})$ lies in $\left(\log\left(\frac{\varepsilon^2}{5m^2}\right), n \log m\right)$, the number of iterations is poly $\left(m, n, \frac{1}{\varepsilon}\right)$. $\square$

Combining the previous lemmas, the following theorem is now immediate by union bounds, where we assume $\varepsilon > 0$ is a small enough constant. This completes the proof of Theorem 3.1.5.

Theorem A.1.6. *With probability at least* $1 - \dfrac{1}{poly(n, m, \frac{1}{\varepsilon})}$, *the local search algorithm in Section 3.4.2 has running time poly*$(n, m, \frac{1}{\varepsilon})$.

Proof. By Lemma A.1.5, if we estimate all the $U_i(\mathbf{x})$ and $\frac{\mathsf{P}U_i(\mathbf{x})}{\mathsf{P}x_j}$ within $\pm\Delta$ additive error, the algorithm ends in poly$(n, m, \frac{1}{\varepsilon})$ iterations. Since within each iteration we estimate U_i and $\frac{\partial U_i(\mathbf{x})}{\partial x_j}$ for poly(m) times, the total number of evaluations is bounded by $N = \text{poly}(n, m, \frac{1}{\varepsilon})$. Then by Lemma A.1.2, we need at most poly($N, m, \frac{1}{\Delta}$) samples for all additive errors to be bounded by $\pm\Delta$ with probability $1 - \frac{1}{\text{poly}(N)}$. Since $\Delta = \frac{\varepsilon^6}{64n \cdot m^5} = \frac{1}{\text{poly}(n,m,\frac{1}{\varepsilon})}$ and the number of iterations is bounded by poly($n, m, \frac{1}{\varepsilon}$), the total running time including sampling steps is poly($n, m, \frac{1}{\varepsilon}$). $\square$

A.2 Improved Approximation for Additive Utilities

In this section, we present a 9.27-utility core for the special case of additive utilities. Recall that for such utilities, $u_i(S) = \sum_{j \in S} u_{ij}$, where $S \subseteq C$ is a subset of candidates and $i \in W$ is a voter. Though we could use the approach as in the submodular-utility case, we lose constant factors first because the local optimum to the Nash Welfare objective only finds an approximate fractional core, and secondly because the randomized rounding needs to scale down budgets to satisfy the size constraints.

We address the first issue by using an exact core solution to the fractional problem. We do this via the classic Lindahl equilibrium that we describe in Section A.2.1. To address the second issue, we use dependent rounding that preserves the budget constraint with probability 1 for additive utilities, and we describe this in Section A.2.2. This yields a 9.27-utility core, though we do not know how to implement the resulting algorithm in polynomial time, since the fractional solution is now via a fixed point argument.

A.2.1 Fractional Solution: Lindahl Equilibrium

Our algorithm for constructing a committee in the approximate core will make use of the Lindahl equilibrium [Lin58, Fol70]. This equilibrium yields a fractional committee that lies in the fractional core (Definition 3.2.2 for $\alpha = 1$).

We follow the approach in [FGM16] for specifying the Lindahl equilibrium. Let

$x_j \geq 0$ denote the fraction to which candidate j is chosen. Here, we will assume for technical reasons that this can be a quantity greater than 1. We assign endowment $\frac{b}{n}$ to each voter $i \in V$, and a price p_{ij} of $j \in C$ for $i \in V$. The Lindahl equilibrium is now defined as follows.

Definition A.2.1 (Lindahl Equilibrium)**.** Let p_{ij} be the price of $j \in C$ for $i \in V$, and let $x_j \geq 0$ be the fraction with which item j is allocated. The prices and allocations constitute a *Lindahl Equilibrium* if:

1. For all $i \in V$, suppose the voter computes allocation $\mathbf{y} \geq 0$ that maximizes her utility $U_i(\mathbf{y})$ subject to her endowment constraint $\sum_{j \in C} p_{ij} y_j \leq \frac{b}{n}$, then we have that $\mathbf{y} = \mathbf{x}$.

2. The profit of the allocation given by $\sum_{i \in V} \sum_{j \in C} y_j p_{ij} - \sum_{j \in C} s_j y_j$, is maximized when $\mathbf{y} = \mathbf{x}$.

The main result in [Fol70] is the following theorem proved by a fixed point argument. (The setting in [Fol70] is much more general.)

Theorem A.2.2 ([Fol70])**.** *If the utility functions $U_i(\mathbf{x})$ have continuous derivatives, are strictly increasing, and are strictly concave, then a Lindahl equilibrium always exists.*

Properties of the Lindahl Equilibrium

We will need the following corollary to Theorem A.2.2.

Corollary A.2.3. *It holds for a Lindahl equilibrium that:*

1. *For all voters $i \in V$, we have $\sum_j p_{ij} x_j = \frac{b}{n}$.*

2. *For each candidate with $x_j > 0$, we have $\sum_i p_{ij} = s_j$, and for each candidate with $x_j = 0$, we have $\sum_i p_{ij} \leq s_j$.*

3. $\sum_{j \in C} s_j x_j = b$.

Given Corollary A.2.3, it is easy and instructive to see that a Lindahl equilibrium lies in the core. Since this proof idea will be crucial to our analysis, we present it for completeness.

Corollary A.2.4 ([Fol70]). *The Lindahl equilibrium lies in the fractional core (Definition 3.2.2 for $\alpha = 1$).*

Proof. Given the equilibrium $\mathbf{x}$, suppose for the purpose of contradiction that there exists a subset $S \subseteq V$ of t voters that can deviate and choose a committee $\mathbf{z}$ of size $\sum_j s_j z_j \leq \frac{t}{n} b$ such that $U_i(\mathbf{z}) > U_i(\mathbf{x})$ for all $i \in V$. Since $\mathbf{x}$ is utility maximizing for the voter i at endowment $\frac{b}{n}$, and since by Item 1 of Corollary A.2.3, this endowment is spent exactly, the committee $\mathbf{z}$ must cost more than b/n. Therefore, for all $i \in V$, we have:

$$\sum_j p_{ij} z_j > \frac{b}{n}.$$

However, by Item 2 of Corollary A.2.3, $\sum_i p_{ij} \leq s_j$ for all $i \in S, j \in C$. Summing the previous inequality over all $i \in S$, and applying $\sum_i p_{ij} \leq s_j$, we obtain:

$$\sum_j s_j z_j \geq \sum_{j, i \in V} p_{ij} z_j \geq \sum_{i \in S} \left(\sum_j p_{ij} z_j \right) > \sum_{i \in S} \frac{b}{n} = t \frac{b}{n}.$$

This contradicts the fact that $\mathbf{z}$ could have been purchased with the endowments of $i \in S$, that is, $\sum_j s_j z_j \leq t \frac{b}{n}$, and thus completes the proof. □

Lindahl Equilibrium for Additive Utilities

So far, we have presented the Lindahl equilibrium in its generality for continuous, concave, non-decreasing utilities. We now specialize it to additive utilities. Given the additive utility function u_i, we make it continuous and unconstrained using the following natural definition:

$$U_i(\mathbf{x}) := \sum_{j \in C} \min(1, x_j) \cdot u_i(\{j\}).$$

It is clearly a concave function.[1] Given this utility function, the Lindahl equilibrium is defined exactly as in Definition A.2.1, and this implies Corollary A.2.3 holds as is. The Lindahl equilibrium computation will be encapsulated by a subroutine below.

Definition A.2.5 (Subroutine $\textsc{Lindahl}(C, W, \{U_i\}, B)$)**.** Given the set of candidates C, a subset $W \subseteq V$ of voters with continuous and concave utilities $\{U_i\}$, and a size constraint B, this procedure finds a fractional solution $\mathbf{x} \geq 0$ over C with $\sum_{j \in C} s_j x_j \leq B$, such that this solution is a Lindahl equilibrium.

Note that the size B could be different from b. Using Definition 3.2.2, the solution $\mathbf{x}$ of this subroutine satisfies the following condition: There is no subset $S \subseteq W$ of voters who can find a committee $\mathbf{z} \geq 0$ over E, with $\sum_{j \in E} s_j z_j \leq \frac{|S|}{|W|} \cdot B$ such that $U_i(\mathbf{z}) > U_i(\mathbf{x})$ for all $i \in S$.

A.2.2 Randomized Rounding and Satisfied Voters

Note that if in the Lindahl equilibrium, we have $x_j > 1$ for some $j \in C$, then all items must be integrally allocated. This is because we assumed $U_i(\mathbf{x}) = u_i^+(\mathbf{y})$ where $y_j = \min(1, x_j)$. This means that if $x_j > 1$, then it can be reduced to 1 without affecting any utilities. Therefore, if some item is fractionally allocated, we can decrease the allocation of j with $x_j > 1$ and increase the allocation of a fractional item, improving the utility of some agent. This violates the core condition. We will therefore assume throughout that $x_j \in [0, 1]$ for all $j \in C$.

[1] As a technicality, to make the function satisfy the preconditions in Theorem A.2.2, we perturb U_i slightly to make it be strictly increasing, strictly concave, and have continuous derivative. This perturbation preserves U_i to within a $(1 + \epsilon)$ approximation for $\epsilon > 0$ being arbitrary small, which will suffice for our approximation guarantees.

Rounding Procedure $\textsc{Round}(C, W, \{u_i\}, B)$

Assume we are given a fractional solution $\mathbf{x} \geq 0$ to $\textsc{Lindahl}(C, W, \{u_{ij}\}_{i,j}, B)$. The procedure $\textsc{Round}(C, W, \{u_{ij}\}, B)$ uses the algorithm of [BPR+17] to the fractionally allocated items in this solution. This produces a committee (O, ℓ), where the candidates in O are chosen integrally, at most candidate ℓ is chosen fractionally, and the total size is always at most B.

Let X_j be the random variable denoting whether candidate j is selected: $X_j = 1$ if $j \in O$, $X_j = 0$ of $j \notin O \cup \{\ell\}$, and X_j is the fractional weight if $j = \ell$. Then the procedure $\textsc{Round}$ preserves the marginals: $\mathbb{E}[X_j] = \min(1, x_j)$ for all $j \in C$, and the $\{X_j\}$ are negatively correlated [GKPS06, PS97]. Note that the candidates with $X_j = 1$ define O, while the one with $X_\ell \in (0, 1)$ is the fractionally chosen candidate (it may not exist).

Note that

$$\mathbb{E}\left[u_{i\ell} + \sum_{j \in O} u_{ij}\right] \geq \mathbb{E}\left[\sum_{j \in C} u_{ij} X_j\right] = \sum_{j \in C} u_{ij} \min(1, x_j) = U_i(\mathbf{x}) \tag{A.2}$$

where the expectation is over the random choice of O. Further, the size constraint is never violated, so that:

$$\sum_{j \in O} s_j \leq \sum_{j \in C} s_j X_j \leq B$$

regardless of the outcome of the rounding procedure. Note finally that any candidate that is fully allocated by $\mathbf{x}$, that is, with $x_j \geq 1$, must be present in O.

Satisfied Voters

The procedure $\textsc{Round}(C, W, \{u_{ij}\}, B)$ returns the committee O. Given the definition of γ-satisfied from Definition 3.4.7, we can strengthen Theorem 3.4.8 as follows:

Theorem A.2.6 (Constant Fraction of Constant-Satisfied Voters)**.** *Given the fractional solution* $\mathbf{x}$ *produced by* $\textsc{Lindahl}(C, W, \{u_{ij}\}, B)$ *where* $|W| = n'$*, there is a integral committee* O *produced by* $\textsc{Round}(C, W, \{u_{ij}\}, B)$ *with at least* $(1-\beta)n'$ γ*-satisfied voters, where* $\beta = \gamma e^{1-\gamma}$*.*

Proving Theorem A.2.6 requires a concentration bound for the sum of negatively correlated weighted Bernoulli random variables.

Lemma A.2.7 ([PS97])**.** *Let* $B_1, B_2, \ldots, B_k$ *be* k *negatively correlated Bernoulli random variables. Let* $X = \sum_{i=1}^{k} \beta_i B_i$*, where* $\beta_i \in [0,1]$*. Let* $\mathbb{E}[X] \geq \mu$*. Then for any constant* $\delta \in (0,1)$ *we have:*

$$\Pr[X < (1-\delta)\mu] < \left(\frac{e^{-\delta}}{(1-\delta)^{(1-\delta)}}\right)^{\mu}.$$

Using the above bound, we complete the proof of Theorem A.2.6.

Proof of Theorem A.2.6. Let $I \subseteq C$ be the set of fully allocated candidates in $\textsc{Lindahl}(C, W, \{u_{ij}\}_{i,j}, B)$ (i.e. $I = \{j \in C \mid x_j \geq 1\}$), and let $F \subseteq C$ be the set of fractionally allocated ones (i.e. $F = \{j \in C \mid x_j \in (0,1)\}$). Let $u_i(F) = \sum_{j \in F} x_j u_{ij}$. There are two cases:

Case (1): Suppose these is some candidate $q \in F$ so that $u_{iq} \geq \frac{u_i(F)}{\gamma}$. In this case, voter i is γ-satisfied with probability 1, since

$$u_{iq} + \sum_{j \in O} u_{ij} \geq \frac{\sum_{j \in F} x_j u_{ij}}{\gamma} + \sum_{j \in I} u_{ij} \geq \frac{\sum_{j \in E} u_{ij} \min(1, x_j)}{\gamma},$$

where the first step uses that $I \subseteq O$.

Case (2): Suppose for every candidate $j \in F$, $u_{ij} < \frac{u_i(F)}{\gamma}$. In this case, we consider the candidate ℓ from the procedure $\textsc{Round}$ as the additament, and invoke our concentration bound of Lemma A.2.7. Notice X_j's are negatively correlated

Bernoulli random variables. Let $\mu = \mathbb{E}\left[\sum_{j\in F} u_{ij} \cdot \frac{\gamma}{u_i(F)} \cdot X_j\right] = \gamma$ and $\delta = 1 - \frac{1}{\gamma}$. By Lemma A.2.7, we have

$$\Pr\left[\sum_{j\in F} u_{ij} \cdot \frac{\gamma}{u_i(F)} \cdot X_j < 1\right] < \left(\frac{e^{-\delta}}{(1-\delta)^{(1-\delta)}}\right)^{\mu} = \gamma e^{1-\gamma}.$$

If the event $\sum_{j\in F} u_{ij} \cdot \frac{\gamma}{u_i(F)} \cdot X_j < 1$ does not happen, the utility of i with additament ℓ is at least

$$\begin{aligned} u_{i\ell} + \sum_{j\in O} u_{ij} &\geq u_{i\ell} X_\ell + \sum_{j\in F\setminus\{\ell\}} u_{ij} X_j + \sum_{j\in I} u_{ij} = \sum_{j\in F} u_{ij} X_j + \sum_{j\in I} u_{ij} \\ &\geq \frac{u_i(F)}{\gamma} + \sum_{j\in I} u_{ij} \geq \frac{\sum_{j\in E} u_{ij} \min(1, x_j)}{\gamma}, \end{aligned}$$

where the first step again uses that $I \subseteq O$. $\square$

A.2.3 The Constant Approximation to the Core

In this section, we modify the algorithm in Section 3.4.4 and prove that it returns a committee in an $O(1)$-utility core. The analysis of our algorithm now critically requires the market-clearing properties of the Lindahl equilibrium presented in Definition A.2.1, Corollary A.2.3, and Corollary A.2.4.

Algorithm. The algorithm is the same as Algorithm 3, except the following lines:

- **Line 7:** $\mathbf{x}_t \leftarrow$ LINDAHL$(C, V_t, \{U_i\}, b_t)$.
- **Line 8:** $O_t \leftarrow$ Solution of ROUND$(C, V_t, \{u_i\}, b_t)$ that satisfies Theorem A.2.6.

Analysis. The analysis follows the same outline as that in Section 3.4.4. First, using the same argument as in that section, the solution T^* is feasible for the size b.

As before, we proceed to show a contradiction. Let S denote the set of voters that deviate, and let A denote the set of items they deviate to. We have $\sum_{j\in A} s_j \leq$

$\frac{|S|}{n} \cdot b$, and A provides an α-factor larger utility to voters in S even after including any additament. Consider the voters in W_t, and let $S_t = S \cap W_t$. These voters are γ-satisfied by O_t with respect to the fractional solution $\mathbf{x}_t$ (Definition 3.4.7). Therefore, if $i \in S_t$ deviates to A to obtain an α-factor larger utility, it must be that $U_i(A) \geq \frac{\alpha}{\gamma} \cdot U_i(\mathbf{x}_t)$. Let $\theta = \frac{\alpha}{\gamma}$. We will assume $\theta > 1$ below.

We will now show the analog of Theorem 3.4.4 using the prices computed by the Lindahl equilibrium. Let p^t denote the prices computed by $\textsc{Lindahl}(C, V_t, \{U_i\}, b_t)$, and let $n_t = |V_t|$. The following lemma generalizes Corollary A.2.4 and bounds the price of the set A via the optimality conditions of the Lindahl equilibrium $\mathbf{x}_t$.

Lemma A.2.8. *For all $i \in S_t$, we have $\sum_{j \in A} p^t_{ij} \geq \theta \cdot \frac{b_t}{n_t}$.*

Proof. First note that $x_{tj} \leq 1$ for all items j. If any $x_{tj} > 1$, then all items must be integrally allocated (see Section A.2.1), which means A cannot achieve a θ-factor larger utility. Note that $\mathbf{x}$ is the utility maximizing solution to a packing problem for voter i where the "size" of item j is p^t_{ij} and the "size" constraint is $\frac{b_t}{n_t}$. This constraint is exactly satisfied by Item 1 of Corollary A.2.3. Since U_i is concave, any solution that produces θ factor more utility must have "size" at least θ times larger. Therefore, the price of A is θ times larger than that of $\mathbf{x}_t$, completing the proof. $\square$

We now bound the size of set S_t as follows:

Lemma A.2.9 (Analog of Theorem 3.4.4)**.** $|S_t| \leq \frac{n_t}{b_t} \cdot \frac{\sum_{j \in A} s_j}{\theta}$.

Proof. Summing the bound in Lemma A.2.8 over $i \in S_t$, we have

$$\sum_{i \in S_t} \sum_{j \in A} p^t_{ij} \geq |S_t| \cdot \theta \cdot \frac{b_t}{n_t}.$$

By Item 2 of Corollary A.2.3, we have

$$\sum_{i \in V_t} p^t_{ij} \leq s_j \quad \forall j \qquad \implies \qquad \sum_{i \in S_t} \sum_{j \in A} p^t_{ij} \leq \sum_{j \in A} s_j.$$

Combining these inequalities completes the proof. $\square$

This proves the analog of Theorem 3.4.4. Continuing as before, we sum Lemma A.2.9 over all t, and using $\sum_{j \in A} s_j \leq \frac{|S|}{n} b$, we have:

$$|S| = \sum_t |S_t| \leq \frac{\sum_{j \in A} s_j}{\theta} \cdot \sum_t \frac{n_t}{b_t} \leq \frac{|S|}{n} \cdot \frac{b}{\theta} \cdot \sum_t \frac{n_t}{b_t}.$$

Therefore, for a blocking coalition to exist, we need:

$$\frac{b}{n} \cdot \sum_t \frac{n_t}{b_t} \geq \theta = \frac{\alpha}{\gamma}. \tag{A.3}$$

We will now set the parameters ω, γ, α so that the above inequality is false. First note by Theorem A.2.6 that $n_{t+1} \leq \beta n_t$ where $\beta = \gamma e^{1-\gamma}$. Further, $b_{t+1} = \omega b_t$. Therefore,

$$\frac{n_{t+1}}{b_{t+1}} \leq \frac{\beta}{\omega} \cdot \frac{n_t}{b_t}$$

with $\frac{n_0}{b_0} = \frac{n}{(1-\omega)b}$. Therefore,

$$\frac{b}{n} \cdot \sum_t \frac{n_t}{b_t} \leq \frac{1}{(1-\omega)} \cdot \sum_{t \geq 0} \left(\frac{\beta}{\omega}\right)^t = \frac{\omega}{(1-\omega)(\omega-\beta)}. \tag{A.4}$$

Combining Eq. (A.3) and Eq. (A.4), for a blocking coalition to exist, we need

$$\alpha \leq \frac{\omega\gamma}{(1-\omega)(\omega-\beta)} = \frac{\omega\gamma}{(1-\omega)(\omega-\gamma e^{1-\gamma})},$$

where the last step uses $\beta = \gamma e^{1-\gamma}$ according to Theorem A.2.6.

For an α slightly larger than the right-hand side, a blocking coalition will therefore not exist. Plugging $\omega = 0.15$ and $\gamma = 6.7$ shows $\alpha < 9.27$, which yields the following theorem:

Theorem 3.1.6 (Additive Utilities). *For additive utilities, a* 9.27*-core is always non-empty.*

A.3 Expected Ranks in a Random Committee

The following is a well-known proof of the statement that if we pick a random size-k subset of $\mathcal{C}$, the t^{th} smallest rank is $t \cdot \frac{m+1}{k+1}$ in expectation. Mark $m+1$ points on a circle. Pick a subset of $k+1$ points uniformly at random, and then choose one point P of these $k+1$ as the cut-off point uniformly at random. Starting from P and going clockwise, mark the next point as the candidate with rank 1, and the point after that as the candidate with rank 2, and so on, until the last point which is marked as the candidate with rank m. The picked subset comprises P and a uniformly random size-k subset of $\mathcal{C}$. By symmetry, the expected clockwise distance going from the t^{th}-smallest ranked chosen candidate to the $(t+1)^{\text{st}}$ is the same for every $t \in \{0, 1, \ldots, k\}$, if we view P as simultaneously the 0^{th} and the $(k+1)^{\text{st}}$ smallest. Since these $k+1$ distances sum to $m+1$, all of them should be $\frac{m+1}{k+1}$.

A.4 Improving the Constant of Theorem 5.1.1

To improve the approximation constant of Theorem 5.1.1, instead of setting $\mu(x)$ to be the median of $F_{|\geq x}$, we introduce a parameter λ to control the quantile of $\mu(x)$ in $F_{|\geq x}$:

$$\mu(x) := F^{(-1)}\Big(\lambda + (1-\lambda)F(x)\Big).$$

Similar to the proof of Theorem 5.1.1, we have

$$\mathsf{FB} = \int_0^1 \int_c^1 (v-c)\,\mathrm{d}F(v)\,\mathrm{d}G(c) \leq \frac{1}{1-\lambda}\int_0^1 \int_{\mu(c)}^1 (v-c)\,\mathrm{d}F(v)\,\mathrm{d}G(c),$$

$$\mathsf{SellerP} \geq \int_0^1 \int_{\mu(c)}^1 \Big(\mu(c) - c\Big)\,\mathrm{d}F(v)\,\mathrm{d}G(c),$$

and

$$\mathsf{BuyerP} \geq \int_{\mu^{(2)}(0)}^{1} \int_{0}^{\mu^{(-2)}(v)} \left(v - \mu^{(-2)}(v)\right) \,\mathrm{d}G(c)\,\mathrm{d}F(v)$$
$$= \int_{0}^{1} \int_{\mu^{(2)}(c)}^{1} \left(v - \mu^{(-2)}(v)\right) \,\mathrm{d}F(v)\,\mathrm{d}G(c).$$

Let $\mathsf{FB}(c) = \frac{1}{1-\lambda}\int_{\mu(c)}^{1}(v-c)\,\mathrm{d}F(v)$, $\mathsf{SellerP}(c) = \int_{\mu(c)}^{1}(\mu(c)-c)\,\mathrm{d}F(v)$ and $\mathsf{BuyerP}(c) = \int_{\mu^{(2)}(c)}^{1}(v - \mu^{(-2)}(v))\,\mathrm{d}F(v)$. We have

$$\mathsf{FB} \leq \int_{0}^{1} \mathsf{FB}(c)\,\mathrm{d}G(c), \quad \mathsf{SellerP} \geq \int_{0}^{1} \mathsf{SellerP}(c)\,\mathrm{d}G(c), \quad \mathsf{BuyerP} \geq \int_{0}^{1} \mathsf{BuyerP}(c)\,\mathrm{d}G(c).$$

Fix any c from now on.

$$\begin{aligned}
\mathsf{FB}(c) &= \frac{1}{1-\lambda}\int_{\mu(c)}^{1} (v-c)\,\mathrm{d}F(v) \\
&= \frac{1}{1-\lambda}\int_{\mu(c)}^{1} \Big(\mu(c) - c\Big)\,\mathrm{d}F(v) + \frac{1}{1-\lambda}\int_{\mu(c)}^{1} \Big(v - \mu(c)\Big)\,\mathrm{d}F(v) \\
&= \frac{1}{1-\lambda}\cdot \mathsf{SellerP}(c) + \frac{1}{1-\lambda}\sum_{k=1}^{\infty}\int_{\mu^{(k)}(c)}^{\mu^{(k+1)}(c)} \Big(v - \mu(c)\Big)\,\mathrm{d}F(v) \\
&\leq \frac{1}{1-\lambda}\cdot \mathsf{SellerP}(c) + \frac{1}{1-\lambda}\sum_{k=1}^{\infty}\int_{\mu^{(k)}(c)}^{\mu^{(k+1)}(c)} \left(\mu^{(k+1)}(c) - \mu(c)\right)\,\mathrm{d}F(v) \\
&= \frac{1}{1-\lambda}\cdot \mathsf{SellerP}(c) + \frac{1}{1-\lambda}\sum_{k=1}^{\infty}\sum_{t=1}^{k}\int_{\mu^{(k)}(c)}^{\mu^{(k+1)}(c)} \left(\mu^{(t+1)}(c) - \mu^{(t)}(c)\right)\,\mathrm{d}F(v) \\
&= \frac{1}{1-\lambda}\cdot \mathsf{SellerP}(c) + \frac{1}{1-\lambda}\sum_{t=1}^{\infty}\sum_{k=t}^{\infty}\int_{\mu^{(k)}(c)}^{\mu^{(k+1)}(c)} \left(\mu^{(t+1)}(c) - \mu^{(t)}(c)\right)\,\mathrm{d}F(v) \\
&= \frac{1}{1-\lambda}\cdot \mathsf{SellerP}(c) + \frac{1}{\lambda(1-\lambda)}\sum_{t=1}^{\infty}\int_{\mu^{(t)}(c)}^{\mu^{(t+1)}(c)} \left(\mu^{(t+1)}(c) - \mu^{(t)}(c)\right)\,\mathrm{d}F(v).
\end{aligned}$$

The last step uses the definition of $\mu(\cdot)$: $k = t$ counts for λ-fraction of the probability of all $k \geq t$.

Additionally, we have

$$
\begin{aligned}
\mathsf{BuyerP}(c) &= \int_{\mu^{(2)}(c)}^{1} \left(v - \mu^{(-2)}(v)\right)\, \mathrm{d}F(v) \\
&= \sum_{t=2}^{\infty} \int_{\mu^{(t)}(c)}^{\mu^{(t+1)}(c)} \left(v - \mu^{(-2)}(v)\right)\, \mathrm{d}F(v) \\
&\geq \sum_{t=2}^{\infty} \int_{\mu^{(t)}(c)}^{\mu^{(t+1)}(c)} \left(\mu^{(t)}(c) - \mu^{(t-1)}(c)\right)\, \mathrm{d}F(v) \\
&= \sum_{t=1}^{\infty} \int_{\mu^{(t+1)}(c)}^{\mu^{(t+2)}(c)} \left(\mu^{(t+1)}(c) - \mu^{(t)}(c)\right)\, \mathrm{d}F(v) \\
&= (1-\lambda) \sum_{t=1}^{\infty} \int_{\mu^{(t)}(c)}^{\mu^{(t+1)}(c)} \left(\mu^{(t+1)}(c) - \mu^{(t)}(c)\right)\, \mathrm{d}F(v).
\end{aligned}
$$

Therefore, $\mathsf{FB}(c) \leq \frac{1}{1-\lambda} \cdot \mathsf{SellerP}(c) + \frac{1}{\lambda(1-\lambda)^2} \cdot \mathsf{BuyerP}(c)$. Setting $\lambda = 0.311$ gives $\mathsf{FB} \leq 8.23 \cdot \max(\mathsf{SellerP}, \mathsf{BuyerP})$.

A.5 Lower Bound Instances

We now prove Theorems 6.1.1 and 6.1.2. We note that though our lower bounds assume that the seller runs interim IR and ex-post IC mechanisms, and the intermediary can send arbitrary signals, the same lower bounds hold even when the seller runs ex-post IC and IR mechanisms, provided we restrict the intermediary's signals to induce posteriors that are product distributions.

Our lower bounds are based on a 2-buyer instance illustrated in Fig. A.1: given values $a > b > c > d$, buyer 1 has value $V_1 \in \{a, b\}$ with probabilities α and $1-\alpha$ respectively, while buyer 2 has value $V_2 \in \{c, d\}$ with probabilities β and $1-\beta$ respectively. We choose $\alpha a = b$ and $\beta c = d$; thus, the virtual values satisfy: $\varphi_1(a) = a$, $\varphi_2(c) = c$, and $\varphi_1(b) = \varphi_2(d) = 0$. Call this distribution $\mathcal{D}$.

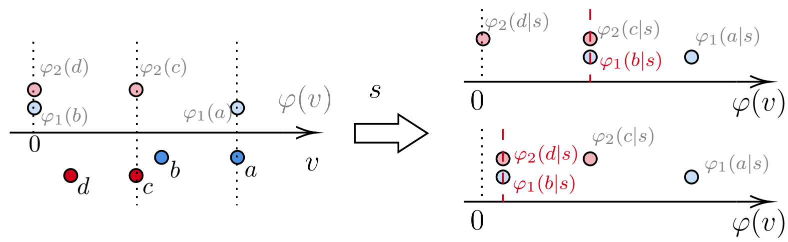

FIGURE A.1: Illustrating the setting for Theorems 6.1.1 and 6.1.2: On the left (below the axis) we show the setting without signaling, where buyer 1 (blue) has values (a, b) and buyer 2 (red) has values (c, d); we also show the corresponding virtual values (above the axis). On the right, we show the two settings characterized by Theorem A.5.2 under which a signal s has non-zero consumer surplus (the changed virtual values are highlighted).

Characterization of optimal signaling. By Theorem 6.2.1, we know any signal that correlates the buyers raises zero consumer surplus. Therefore, the only signals $\mathcal{S}$ of interest are those under which buyer values are independent. Abusing notation we denote such a signal as $s = (\alpha', \beta')$, where $\Pr[v_1 = a] = \alpha'$ and $\Pr[v_2 = c] = \beta'$. Note that in this instance, for a signal to get maximum welfare the resulting optimal mechanism must always award buyer 1, and for non-zero consumer surplus it must award the item to buyer 1 at price b, or buyer 2 at price d.

Let $\mathrm{CS}(s)$ denote the consumer surplus under any such a signal s, and $\varphi_1(b|s)$ and $\varphi_2(d|s)$ denote the new virtual values (note that by definition, $\varphi_1(a|s) = a$ and $\varphi_2(c|s) = c$ under any signal s with $\alpha', \beta' > 0$). We can use Myerson's characterization (Section 6.2.1) to exhaustively characterize the resulting optimal mechanisms as a function of $(\varphi_1(b|s), \varphi_2(d|s))$:

Proposition A.5.1. *Conditioned on receiving a signal s, we have the following cases:*

(1) If $\varphi_1(b) \geq c$, then the optimal mechanism is to sell to Buyer 1 at price b. $\mathrm{CS}(s) = (a - b)\alpha'$.

(2) If $\varphi_2(d) \geq \max(0, \varphi_1(b))$, then the optimal mechanism is to try selling to Buyer 1 at price a then to Buyer 2 at price d. $\mathrm{CS}(s) = (1 - \alpha')\beta'(c - d)$.

(3) If $\varphi_1(b) \leq 0$ and $\varphi_2(d) \leq 0$, then the optimal mechanism is to try selling to Buyer 1 at price a then to Buyer 2 at price c. $\mathrm{CS}(s) = 0$.

(4) If $0 \leq \varphi_1(b) \leq c$ and $\varphi_2(d) \leq \varphi_1(b)$, then the optimal mechanism is to sell to Buyer 1 at price b if Buyer 2 has valuation d; otherwise, it tries selling to Buyer 1 at price a then to Buyer 2 at price c. $\mathrm{CS}(s) = \alpha'(1-\beta')(a-b)$.

Our main insight, however, is that the setting can be further simplified to get the following structural property for the optimal signaling scheme.

Theorem A.5.2 (Structural Theorem)**.** *In an optimal signaling scheme, the only signals $s = (\alpha', \beta')$ that raise non-zero consumer surplus have the following form:*

(1') Under signal s, $\varphi_1(b|s) = \varphi_2(c|s) = c$ and $\mathrm{CS}(s) = \alpha'(a-b)$.

(2') Under signal s, $\varphi_2(d|s) = \varphi_1(b|s) \geq 0$ and $\mathrm{CS}(s) = \alpha'(1-\beta')(a-b)$.

Proof. Recall that we restrict ourselves to signals $\mathcal{S}$ under which the buyer valuations remain independent. Any such signal can be alternately written as $s = (\alpha', \beta')$ where $\alpha' = \Pr[v_1 = a]$ and $\beta' = \Pr[v_2 = c]$. For ease of notation, we henceforth drop the conditioning of virtual valuations on signal s (i.e., write $\varphi(\cdot)$ for $\varphi(\cdot|s)$) when clear from context.

Next, let γ_s denote the weight of any signal $s = (\alpha', \beta')$. The signaling scheme that maximizes consumer surplus is the solution to the following linear program written over signals $s = (\alpha', \beta')$:

$$
\begin{array}{lrll}
\text{Maximize} & \displaystyle\sum_s \gamma_s \mathrm{CS}(s) & & \\
\text{Subject to} & \sum_{s=(\alpha',\beta')} \gamma_s \alpha'\beta' & \leq \alpha\beta & \\
& \sum_{s=(\alpha',\beta')} \gamma_s \alpha'(1-\beta') & \leq \alpha(1-\beta) & \\
& \sum_{s=(\alpha',\beta')} \gamma_s (1-\alpha')\beta' & \leq (1-\alpha)\beta & \\
& \sum_{s=(\alpha',\beta')} \gamma_s (1-\alpha')(1-\beta') & \leq (1-\alpha)(1-\beta) & \\
& \gamma_s & \geq 0 & \forall s
\end{array}
\tag{A.5}
$$

We examine the cases in Proposition A.5.1 with positive consumer surplus, and characterize the optimal solution:

- In Case (1), we have $\varphi_1(b) = c$. To see this, consider any signal s with $\varphi_1(b) > c$. Suppose we increase α' and decrease γ_s while preserving the product $\alpha'\gamma_s$. Since $\gamma_s \mathrm{CS}(s) = \gamma_s \alpha'(a-b)$, this is preserved by the change. Therefore, the objective of LP (A.5) is preserved, and so are the first two constraints. Further, since $(1-\alpha')$ and γ_s decrease, this only makes the third and fourth constraints more feasible. This transformation decreases $\varphi_1(b)$.

- In Case (2) and (4), we have $\varphi_2(d) = \varphi_1(b)$. It does not help to make them unequal by a similar argument as above: In case (2), if $\varphi_2(d) > \varphi_1(b)$, we can increase β' while preserving $\gamma_s\beta'$. Since $\gamma_s \mathrm{CS}(s) = \gamma_s(1-\alpha')\beta'(c-d)$, this does not change the contribution to the objective of LP (A.5), and preserves all constraints. This transformation decreases $\varphi_2(d)$. In case (4), if $\varphi_2(d) < \varphi_1(b)$, we can increase α' while preserving $\gamma_s\alpha'$. Since $\gamma_s \mathrm{CS}(s) = \gamma_s\alpha'(1-\beta')(a-b)$, this does not change the contribution to the objective of LP (A.5), and preserves all constraints. This transformation decreases $\varphi_1(b)$.

Therefore, the only two types of signals s that give positive CS are

(1') If $\varphi_1(b) = c$, then $\mathrm{CS}(s) = \alpha'(a-b)$.

(2') If $\varphi_2(d) = \varphi_1(b) \geq 0$, then $\mathrm{CS}(s) = \max((1-\alpha')\beta'(c-d), \alpha'(1-\beta')(a-b))$.

As $(1-\beta')(b-d) \geq 0$, we have

$$\left(b - d + \frac{\beta'}{1-\beta'}(c-d)\right)(1-\beta') \geq \beta'(c-d).$$

Notice that in Case (2'), we have $\varphi_1(b) = b - \frac{\alpha'}{1-\alpha'}(a-b) = d - \frac{\beta'}{1-\beta'}(c-d) = \varphi_2(d)$. This gives

$$\alpha'(1-\beta')(a-b) \geq (1-\alpha')\beta'(c-d).$$

Thus, the two types of signals s that give positive CS become

(1') If $\varphi_1(b) = c$, then $\text{CS}(s) = \alpha'(a-b)$.

(2') If $\varphi_2(d) = \varphi_1(b) \geq 0$, then $\text{CS}(s) = \alpha'(1-\beta')(a-b)$. $\square$

Using the above structural theorem, the proofs of Theorems 6.1.1 and 6.1.2 follow by different choices of the parameters (a, b, c, d). Suppose the virtual values of b and d are slightly above zero with $\varphi_1(b) > \varphi_2(d)$ so that Case (4) in Proposition A.5.1 is uniquely optimal for the seller. The optimal auction generates consumer surplus $\text{CS}(\mathcal{D}) = \alpha(1-\beta)(a-b) = \frac{b}{c} \cdot \frac{c-d}{a} \cdot (a-b)$ according to Proposition A.5.1.

A.5.1 Proof of Theorem 6.1.1

To prove Theorem 6.1.1, we set $b \to c^+$.[2] Now in Proposition A.5.1, in Case (1), we must have $\alpha' \to 0^+$ since $b \to c^+$, so that $\text{CS} \to 0$. Also if $\alpha' = 1$ in a signal then $\text{CS} = 0$ here. The only other signal where the item is allocated to the higher bidder is in Case (4) when $\beta' = 0$. Let γ denote the probability of the signal of this type $s = (\alpha', 0)$. (Having multiple signals of this form gives the same CS as having a single signal as their average.) Since $\varphi_1(b) \geq \varphi_2(d)$, we have $\alpha' \leq \frac{b-d}{a-d}$.

By the constraints of LP (A.5), we have:

$$\Pr[v_1 = b \wedge v_2 = d] = \gamma(1-\alpha') \leq (1-\alpha)(1-\beta),$$

which simplifies to $\gamma \leq \frac{(a-d)\cdot(c-d)}{ac}$. The consumer surplus in this case is therefore:

$$\text{CS} = \gamma \text{CS}(s) = \gamma\alpha'(a-b) \leq \frac{(b-d)\cdot(c-d)}{ac} \cdot (a-b) \leq \frac{b-d}{b} \cdot \text{CS}(\mathcal{D}).$$

Setting $d = \left(1 - \frac{\varepsilon}{2}\right) b$ and combining with the fact that $\text{CS} \to 0$ in Case (1), we have the consumer surplus of any efficient signaling, $\text{CS} \to \frac{\varepsilon}{2} \cdot \text{CS}(\mathcal{D})$ so $\text{CS} < \varepsilon \cdot \text{CS}(\mathcal{D})$.

[2] $\varphi_1(b) - \varphi_2(d)$ and $\varphi_2(d)$ can be arbitrarily small as long as positive, so we take the limits for them first, i.e., we are calculating $\lim_{b \to c^+} \lim_{\varphi_2(d) \to 0^+, \varphi_1(b) \to \varphi_2(d)^+} \text{CS}$ in the following part of the proof. This allows us to treat $\alpha = \frac{b}{a}$ and $\beta = \frac{d}{c}$ in calculating $(1-\alpha)(1-\beta)$, as α and β are not infinitesimally close to 1 for any fixed ε. (We will define $d = (1 - \varepsilon/2)b$.)

A.5.2 Proof of Theorem 6.1.2

Without signaling, $\mathbf{E}[\max v_i] = \alpha a + (1-\alpha)b$ and $\mathcal{R}(\mathcal{D}) = \alpha a + (1-\alpha)\beta c$. (Case (2), (3) and (4) in Proposition A.5.1 give the same revenue $\mathcal{R}(\mathcal{D})$.) Therefore

$$\textsc{Opt} = \mathbf{E}[\max v_i] - \mathcal{R}(\mathcal{D}) = (1-\alpha)(b - \beta c).$$

Now we assign the values as: $\alpha = \beta = 1-\delta$; $a = \frac{1}{(1-\delta)^2}$, $b = \frac{1}{1-\delta}$, $c = 1$, $d = 1-\delta$ with $\delta \to 0^+$. Then we plug in the values and the two possible types of signals s in Theorem A.5.2 become

(1') If $\alpha' = \frac{1-\delta}{2-\delta} < \frac{1}{2}$, then $\mathrm{CS}(s) \le \frac{1}{2}\delta(1+o(1))$.

(2') If $\alpha' \le 1-\delta$; $\beta' = \frac{1-3(1-\alpha')+3\delta(1-\alpha')-\delta^2(1-\alpha')}{1-2(1-\alpha')+\delta(1-\alpha')} > \frac{1-3(1-\alpha')}{1-2(1-\alpha')}$, then $\mathrm{CS}(s) \le \alpha'(1-\beta')\delta(1+o(1))$.

The consumer surplus maximizing signaling scheme should use t signals $S_{2,i}$ of type (2'), with $\alpha'_{2,i}$, $\beta'_{2,i}$ and weight $w(S_{2,i})$. There is an additional signal S_1 (with weight $w(S_1)$) of type (1') with α'_1 and β'_1. (Having multiple signals of type (1') gives the same CS as having a single signal as their average.) Denoting the valuation of the first buyer by v_1 and the second buyer by v_2, the constraints in LP (A.5) imply the two constraints:

$$\Pr[v_1 = b] = (1-\alpha'_1)w(S_1) + \sum_{i=1}^{t}(1-\alpha'_{2,i}) \cdot w(S_{2,i}) \le 1-\alpha = \delta, \quad \text{(Constraint (I))}$$

$$\Pr[v_1 = b \wedge v_2 = d] = \sum_{i=1}^{t}(1-\alpha'_{2,i})(1-\beta'_{2,i}) \cdot w(S_{2,i}) \le (1-\alpha)(1-\beta) = \delta^2. \quad \text{(Constraint (II))}$$

Note that $\textsc{Opt} = (1-\alpha)(b-\beta c) = 2\delta^2(1+o(1))$.

The total consumer surplus therefore is:

$$
\begin{aligned}
\text{CS} \leq & \frac{1}{2}\delta(1+o(1)) \cdot w(S_1) + \sum_{i=1}^{t} \alpha'_{2,i}(1-\beta'_{2,i})\delta(1+o(1)) \cdot w(S_{2,i}) \\
\leq & \frac{1}{2}\delta(1+o(1)) \cdot 2\left(\delta - \sum_{i=1}^{t}(1-\alpha'_{2,i}) \cdot w(S_{2,i})\right) + \delta(1+o(1)) \sum_{i=1}^{t} \alpha'_{2,i}(1-\beta'_{2,i}) \cdot w(S_{2,i}) \\
= & \delta^2(1+o(1)) + \delta(1+o(1)) \sum_{i=1}^{t}(\alpha'_{2,i}(1-\beta'_{2,i}) - (1-\alpha'_{2,i})) \cdot w(S_{2,i}) \\
\leq & \delta^2(1+o(1)) + \delta(1+o(1)) \sum_{i=1}^{t}(1-\alpha'_{2,i})(1-\beta'_{2,i}) \cdot w(S_{2,i}) \\
\leq & \delta^2(1+o(1)) + \delta(1+o(1)) \cdot \delta^2 = \delta^2(1+o(1)).
\end{aligned}
$$

Here the second inequality follows from Constraint (I), and $\alpha'_1 < \frac{1}{2}$ by the condition of (1'). The third inequality uses the implication of $\varphi_2(d) = \varphi_1(b)$ that $\beta'_{2,i} > \frac{1-3(1-\alpha'_{2,i})}{1-2(1-\alpha'_{2,i})}$. The fourth inequality uses Constraint (II). This establishes a lower bound of 2, since $\text{OPT} = 2\delta^2(1+o(1))$.

Bibliography

[ABC+17] Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh. Justified representation in approval-based committee voting. *Social Choice and Welfare*, 48(2):461–485, 2017.

[ABE+18] Elliot Anshelevich, Onkar Bhardwaj, Edith Elkind, John Postl, and Piotr Skowron. Approximating optimal social choice under metric preferences. *Artificial Intelligence*, 264:27–51, 2018.

[ABES19] Haris Aziz, Felix Brandt, Edith Elkind, and Piotr Skowron. *Computational Social Choice: The First Ten Years and Beyond*, pages 48–65. Springer International Publishing, 2019.

[ABG+20] Reza Alijani, Siddhartha Banerjee, Sreenivas Gollapudi, Kamesh Munagala, and Kangning Wang. Predict and match: Prophet inequalities with uncertain supply. *Proc. ACM Meas. Anal. Comput. Syst.*, 4(1):04:1–04:23, 2020.

[ABMW22] Reza Alijani, Siddhartha Banerjee, Kamesh Munagala, and Kangning Wang. The limits of an information intermediary in auction design. In *Proceedings of the 2022 ACM Conference on Economics and Computation, EC*, 2022.

[AD54] Kenneth J. Arrow and Gerard Debreu. Existence of an equilibrium for a competitive economy. *Econometrica*, 22(3):265–290, 1954.

[ADSY10] Shipra Agrawal, Yichuan Ding, Amin Saberi, and Yinyu Ye. Correlation robust stochastic optimization. In *SODA*, page 1087–1096, 2010.

[AEH+18] Haris Aziz, Edith Elkind, Shenwei Huang, Martin Lackner, Luis Sánchez Fernández, and Piotr Skowron. On the complexity of extended and proportional justified representation. In *AAAI*, pages 902–909, 2018.

[AFSV21] Elliot Anshelevich, Aris Filos-Ratsikas, Nisarg Shah, and Alexandros A. Voudouris. Distortion in social choice problems: The first 15 years and beyond. In Zhi-Hua Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, pages 4294–4301. ijcai.org, 2021.

[AGK+04] Vijay Arya, Naveen Garg, Rohit Khandekar, Adam Meyerson, Kamesh Munagala, and Vinayaka Pandit. Local search heuristics for k-median and facility location problems. *SIAM J. Comput.*, 33(3):544–562, 2004.

[AP17] Elliot Anshelevich and John Postl. Randomized social choice functions under metric preferences. *Journal of Artificial Intelligence Research*, 58:797–827, 2017.

[AS16a] Elliot Anshelevich and Shreyas Sekar. Blind, greedy, and random: Algorithms for matching and clustering using only ordinal information. In Dale Schuurmans and Michael P. Wellman, editors, *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*, pages 390–396. AAAI Press, 2016.

[AS16b] Elliot Anshelevich and Shreyas Sekar. Truthful mechanisms for matching and clustering in an ordinal world. In Yang Cai and Adrian Vetta, editors, *Web and Internet Economics - 12th International Conference, WINE 2016, Montreal, Canada, December 11-14, 2016, Proceedings*, volume 10123 of *Lecture Notes in Computer Science*, pages 265–278. Springer, 2016.

[AS21] Haris Aziz and Nisarg Shah. Participatory budgeting: models and approaches. In *Pathways Between Social Science and Computational Social Science*, pages 215–236. 2021.

[Ban65] J. Banzhaf. Weighted voting doesn't work: A mathematical analysis. *Rutgers Law Review*, 19(2):317–343, 1965.

[BBM15] Dirk Bergemann, Benjamin Brooks, and Stephen Morris. The limits of price discrimination. *American Economic Review*, 105(3):921–57, 2015.

[BBM17] Dirk Bergemann, Benjamin Brooks, and Stephen Morris. First-price auctions with general information structures: Implications for bidding and revenue. *Econometrica*, 85(1):107–143, 2017.

[BCE+16] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia. *Handbook of computational social choice*. Cambridge University Press, 2016.

[BCGZ18] Moshe Babaioff, Yang Cai, Yannai A. Gonczarowski, and Mingfei Zhao. The best of both worlds: Asymptotically efficient mechanisms with a guarantee on the expected gains-from-trade. In *EC*, page 373, 2018.

[BCH+15] Craig Boutilier, Ioannis Caragiannis, Simi Haber, Tyler Lu, Ariel D Procaccia, and Or Sheffet. Optimal social choice functions: A utilitarian view. *Artificial Intelligence*, 227:190–213, 2015.

[BCWZ17] Johannes Brustle, Yang Cai, Fa Wu, and Mingfei Zhao. Approximating gains from trade in two-sided markets via simple mechanisms. In *EC*, pages 589–590, 2017.

[BD16] Liad Blumrosen and Shahar Dobzinski. (Almost) efficient mechanisms for bilateral trading. *CoRR*, abs/1604.04876, 2016.

[BFST19] Markus Brill, Piotr Faliszewski, Frank Sommer, and Nimrod Talmon. Approximation algorithms for balancedCC multiwinner rules. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '19, Montreal, QC, Canada, May 13-17, 2019*, pages 494–502. International Foundation for Autonomous Agents and Multiagent Systems, 2019.

[BGG20] Moshe Babaioff, Kira Goldner, and Yannai A. Gonczarowski. Bulow-klemperer-style results for welfare maximization in two-sided markets. In *SODA*, pages 2452–2471, 2020.

[BKS07] Steven J. Brams, D. Marc Kilgour, and M. Remzi Sanver. A minimax procedure for electing committees. *Public Choice*, 132(3):401–420, 2007.

[BKV18] Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. Finding fair and efficient allocations. In *EC*, page 557–574, 2018.

[BM16] Liad Blumrosen and Yehonatan Mizrahi. Approximating gains-from-trade in bilateral trading. In *WINE*, pages 400–413, 2016.

[BM19] Dirk Bergemann and Stephen Morris. Information design: A unified perspective. *Journal of Economic Literature*, 57(1):44–95, 2019.

[BMPLZ19] Santiago R. Balseiro, Vahab S. Mirrokni, Renato Paes Leme, and Song Zuo. Dynamic double auctions: towards first best. In *SODA*, pages 157–172, 2019.

[BP07] Dirk Bergemann and Martin Pesendorfer. Information structures in optimal auctions. *Journal of Economic Theory*, 137(1):580–609, 2007.

[BPR+17] Jaroslaw Byrka, Thomas W. Pensyl, Bartosz Rybicki, Aravind Srinivasan, and Khoa Trinh. An improved approximation for k-median and positive correlation in budgeted optimization. *ACM Trans. Algorithms*, 13(2):23:1–23:31, 2017.

[BS05] William C. Brainard and Herbert E. Scarf. How to compute equilibrium prices in 1891. *American Journal of Economics and Sociology*, 64(1):57–83, 2005.

[BSS18] Jaroslaw Byrka, Piotr Skowron, and Krzysztof Sornat. Proportional approval voting, harmonic k-median, and negative association. In *45th International Colloquium on Automata, Languages, and Programming, ICALP*, volume 107 of *LIPIcs*, pages 26:1–26:14, 2018.

[Cab04] Yves Cabannes. Participatory budgeting: a significant contribution to participatory democracy. *Environment and urbanization*, 16(1):27–46, 2004.

[CC83] John R. Chamberlin and Paul N. Courant. Representative deliberations and representative decisions: Proportional representation and the borda rule. *The American Political Science Review*, 77(3):718–733, 1983.

[CCPV11] Gruia Călinescu, Chandra Chekuri, Martin Pál, and Jan Vondrák. Maximizing a monotone submodular function subject to a matroid constraint. *SIAM J. Comput.*, 40(6):1740–1766, 2011.

[CDHW20] Rachel Cummings, Nikhil R Devanur, Zhiyi Huang, and Xiangning Wang. Algorithmic price discrimination. In *Proceedings of the 31st Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 2432–2451. SIAM, 2020.

[CdKLT16] Riccardo Colini-Baldeschi, Bart de Keijzer, Stefano Leonardi, and Stefano Turchetta. Approximately efficient double auctions with strong budget balance. In *SODA*, pages 1424–1443, 2016.

[CEF+20] Yang Cai, Federico Echenique, Hu Fu, Katrina Ligett, Adam Wierman, and Juba Ziani. Third-party data providers ruin simple mechanisms. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 4(1):1–31, 2020.

[CFF+16] Ioannis Caragiannis, Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, Kristoffer Arnsfelt Hansen, and Zihan Tan. Truthful facility assignment with resource augmentation: An exact analysis of serial dictatorship. In Yang Cai and Adrian Vetta, editors, *Web and Internet Economics - 12th International Conference, WINE 2016, Montreal, Canada, December 11-14, 2016, Proceedings*, volume 10123 of *Lecture Notes in Computer Science*, pages 236–250. Springer, 2016.

[CFLM19] Xingyu Chen, Brandon Fain, Liang Lyu, and Kamesh Munagala. Proportionally fair clustering. In *ICML*, pages 1032–1041, 2019.

[CG15] Richard Cole and Vasilis Gkatzelis. Approximating the nash social welfare with indivisible items. *SIGecom Exch.*, 14(1):84–88, November 2015.

[CGdK+17] Riccardo Colini-Baldeschi, Paul W. Goldberg, Bart de Keijzer, Stefano Leonardi, Tim Roughgarden, and Stefano Turchetta. Approximately efficient two-sided combinatorial auctions. In *EC*, pages 591–608, 2017.

[CGMW21] Yu Cheng, Nick Gravin, Kamesh Munagala, and Kangning Wang. A simple mechanism for a budget-constrained buyer. *ACM Trans. Economics and Comput.*, 9(2):10:1–10:25, 2021.

[CGMZ21] Yang Cai, Kira Goldner, Steven Ma, and Mingfei Zhao. On multi-dimensional gains from trade maximization. In *SODA*, pages 1079–1098, 2021.

[CH14] Archishman Chakraborty and Rick Harbaugh. Persuasive puffery. *Marketing Science*, 33(3):382–400, 2014.

[CJMW20] Yu Cheng, Zhihao Jiang, Kamesh Munagala, and Kangning Wang. Group fairness in committee selection. *ACM Trans. Economics and Comput.*, 8(4):23:1–23:18, 2020.

[CM88] Jacques Crémer and Richard P. McLean. Full extraction of the surplus in Bayesian and dominant strategy auctions. *Econometrica*, 56(6):1247–1257, 1988.

[CP11] Ioannis Caragiannis and Ariel D Procaccia. Voting almost maximizes social welfare despite limited communication. *Artificial Intelligence*, 175(9):1655, 2011.

[CR22] Moses Charikar and Prasanna Ramakrishnan. Metric distortion bounds for randomized social choice. In Joseph (Seffi) Naor and Niv Buchbinder, editors, *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9 - 12, 2022*, pages 2986–3004. SIAM, 2022.

[CVZ10] Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Dependent randomized rounding via exchange properties of combinatorial structures. In *FOCS*, pages 575–584, 2010.

[CVZ14] Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Submodular function maximization via the multilinear relaxation and contention resolution schemes. *SIAM J. Comput.*, 43(6):1831–1879, 2014.

[DFL+21] Paul Dütting, Federico Fusco, Philip Lazos, Stefano Leonardi, and Rebecca Reiffenhäuser. Efficient two-sided markets with limited information. In *STOC*, pages 1452–1465, 2021.

[DH66] Otto Anderson Davis and Melvin J Hinich. *Some results related to a mathematical model of policy formation in a democratic society.* Carnegie Institute of Technology, Graduate School of Industrial Administration, 1966.

[DKQ16] Shaddin Dughmi, David Kempe, and Ruixin Qiang. Persuasion with limited communication. In *Proceedings of the 17th ACM Conference on Economics and Computation, EC*, pages 663–680, 2016.

[DMSW21] Yuan Deng, Jieming Mao, Balasubramanian Sivan, and Kangning Wang. Optimal pricing schemes for an impatient buyer. *CoRR*, abs/2106.02149, 2021.

[DMSW22] Yuan Deng, Jieming Mao, Balasubramanian Sivan, and Kangning Wang. Approximately efficient bilateral trade. In *STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20 - 24, 2022*, pages 718–721, 2022.

[Dro81] H. R. Droop. On methods of electing representatives. *Journal of the Statistical Society of London*, 44(2):141–202, 1881.

[DS79] Pradeep Dubey and Lloyd S. Shapley. Mathematical properties of the banzhaf power index. *Mathematics of Operations Research*, 4(2):99–131, 1979.

[Dug17] Shaddin Dughmi. Algorithmic information structure design: A survey. *ACM SIGecom Exchanges*, 15(2):2–24, 2017.

[DX19] Shaddin Dughmi and Haifeng Xu. Algorithmic bayesian persuasion. *SIAM Journal on Computing*, pages 16–68, 2019.

[EFSS17] Edith Elkind, Piotr Faliszewski, Piotr Skowron, and Arkadii Slinko. Properties of multiwinner voting rules. *Social Choice and Welfare*, 48(3):599–632, 2017.

[EG59] Edmund Eisenberg and David Gale. Consensus of subjective probabilities: The pari-mutuel method. *The Annals of Mathematical Statistics*, 30(1):165–168, 1959.

[EH84] James M Enelow and Melvin J Hinich. *The spatial theory of voting: An introduction*. CUP Archive, 1984.

[EH90] James M Enelow and Melvin J Hinich. *Advances in the spatial theory of voting*. Cambridge University Press, 1990.

[Elk07] Edith Elkind. Designing and learning optimal finite support auctions. In *Proceedings of the 18^{th} Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 736–745. SIAM, 2007.

[End17] Ulle Endriss. *Trends in Computational Social Choice*. Lulu.com, 2017.

[Fei98] Uriel Feige. A threshold of $\ln n$ for approximating set cover. *J. ACM*, 45(4):634–652, July 1998.

[FEL$^+$17] Luis Sánchez Fernández, Edith Elkind, Martin Lackner, Norberto Fernández García, Jesús Arias-Fisteus, Pablo Basanta-Val, and Piotr Skowron. Proportional justified representation. In *AAAI*, pages 670–676, 2017.

[FFG16] Michal Feldman, Amos Fiat, and Iddan Golomb. On voting and facility location. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, pages 269–286. ACM, 2016.

[FFZ14] Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, and Jie Zhang. Social welfare in one-sided matchings: Random priority and beyond. In Ron Lavi, editor, *Algorithmic Game Theory - 7th International Symposium, SAGT 2014, Haifa, Israel, September 30 - October 2, 2014. Proceedings*, volume 8768 of *Lecture Notes in Computer Science*, pages 1–12. Springer, 2014.

[FGM16] Brandon Fain, Ashish Goel, and Kamesh Munagala. The core of the participatory budgeting problem. In *WINE*, pages 384–399, 2016.

[FGPS19] Eric J. Friedman, Vasilis Gkatzelis, Christos-Alexandros Psomas, and Scott Shenker. Fair and efficient memory sharing: Confronting free riders. In *AAAI*, pages 1965–1972, 2019.

[Fis77] Peter C Fishburn. Condorcet social choice functions. *SIAM Journal on applied Mathematics*, 33(3):469–489, 1977.

[FLLT18] Hu Fu, Christopher Liaw, Pinyan Lu, and Zhihao Gavin Tang. The value of information concealment. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, page 2533–2544, 2018.

[FLPT18] Piotr Faliszewski, Martin Lackner, Dominik Peters, and Nimrod Talmon. Effective heuristics for committee scoring rules. *Proceedings of the AAAI Conference on Artificial Intelligence, AAAI*, 32(1), Apr. 2018.

[FMS18] Brandon Fain, Kamesh Munagala, and Nisarg Shah. Fair allocation of indivisible public goods. In *EC*, pages 575–592, 2018.

[Fol70] Duncan K Foley. Lindahl's solution and the core of an economy with public goods. *Econometrica*, pages 66–72, 1970.

[GHS20] Vasilis Gkatzelis, Daniel Halpern, and Nisarg Shah. Resolving the optimal metric distortion conjecture. In Sandy Irani, editor, *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 1427–1438. IEEE, 2020.

[Gib73] Allan Gibbard. Manipulation of voting schemes: a general result. *Econometrica: journal of the Econometric Society*, pages 587–601, 1973.

[GKM17] Ashish Goel, Anilesh K Krishnaswamy, and Kamesh Munagala. Metric distortion of social choice rules: Lower bounds and fairness properties. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, pages 287–304. ACM, 2017.

[GKPS06] Rajiv Gandhi, Samir Khuller, Srinivasan Parthasarathy, and Aravind Srinivasan. Dependent rounding and its applications to approximation algorithms. *J. ACM*, 53(3):324–360, 2006.

[GKSA19] Ashish Goel, Anilesh K. Krishnaswamy, Sukolsak Sakshuwong, and Tanja Aitamurto. Knapsack voting for participatory budgeting. *ACM Trans. Econ. Comput.*, 7(2), Jul 2019.

[GMWW22] Guru Guruganesh, Aranyak Mehta, Di Wang, and Kangning Wang. Prior-independent auctions for heterogeneous bidders. *CoRR*, abs/2207.09429, 2022.

[GTW21] Nick Gravin, Zhihao Gavin Tang, and Kangning Wang. Online stochastic matching with edge arrivals. In *48th International Colloquium on Automata, Languages, and Programming, ICALP 2021, July 12-16, 2021, Glasgow, Scotland (Virtual Conference)*, volume 198 of *LIPIcs*, pages 74:1–74:20, 2021.

[HS22] Nima Haghpanah and R. Siegel. The limits of multi-product price discrimination. *American Economic Review*, 2022.

[JMW20] Zhihao Jiang, Kamesh Munagala, and Kangning Wang. Approximately stable committee selection. In *Proccedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020*, pages 463–472, 2020.

[KFMB17] Mayuresh Kunjir, Brandon Fain, Kamesh Munagala, and Shivnath Babu. ROBUS: fair cache allocation for data-parallel workloads. In *SIGMOD*, pages 219–234, 2017.

[KG11] Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.

[KJW$^+$21] Anilesh Kollagunta Krishnaswamy, Zhihao Jiang, Kangning Wang, Yu Cheng, and Kamesh Munagala. Fair for all: Best-effort fairness guarantees for classification. In *The 24th International Conference on Artificial Intelligence and Statistics, AISTATS 2021, April 13-15,* *2021, Virtual Event*, volume 130 of *Proceedings of Machine Learning Research*, pages 3259–3267, 2021.

[KK22] Fatih Erdem Kizilkaya and David Kempe. Plurality veto: A simple voting rule achieving optimal metric distortion. *CoRR*, abs/2206.07098, 2022.

[KM22] Shao-Heng Ko and Kamesh Munagala. Optimal price discrimination for randomized mechanisms. In *Proceedings of the 23rd ACM Conference on Economics and Computation, EC*, 2022.

[KPV22] Zi Yang Kang, Francisco Pernice, and Jan Vondrák. Fixed-price approximations in bilateral trade. In *SODA*, pages 2964–2985, 2022.

[LB11] Tyler Lu and Craig Boutilier. Budgeted social choice: From consensus to personalized decision making. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence, IJCAI*, page 280–286, 2011.

[LFMW21] Liang Lyu, Brandon Fain, Kamesh Munagala, and Kangning Wang. Centrality with diversity. In *WSDM '21, The Fourteenth ACM International Conference on Web Search and Data Mining, Virtual Event, Israel, March 8-12, 2021*, pages 644–652. ACM, 2021.

[Lin58] Erik Lindahl. Just taxation—a positive solution. In *Classics in the theory of public finance*, pages 168–176. 1958.

[McA08] R Preston McAfee. The gains from trade under fixed price mechanisms. *Applied economics research bulletin*, 1(1):1–10, 2008.

[MIMG99] Samuel Merrill III, Samuel Merrill, and Bernard Grofman. *A unified theory of voting: Directional and proximity spatial models.* Cambridge University Press, 1999.

[MLW22] Jieming Mao, Renato Paes Leme, and Kangning Wang. Interactive communication in bilateral trade. In *13th Innovations in Theoretical Computer Science Conference, ITCS 2022, January 31 - February 3, 2022, Berkeley, CA, USA*, volume 215 of *LIPIcs*, pages 105:1–105:21, 2022.

[MMW22] Mohammad Mahdian, Jieming Mao, and Kangning Wang. Regret minimization with noisy observations. *CoRR*, abs/2207.09435, 2022.

[Mon95] Burt L. Monroe. Fully proportional representation. *The American Political Science Review*, 89(4):925–940, 1995.

[Mou86] Hervé Moulin. Choosing from a tournament. *Social Choice and Welfare*, 3(4):271–291, 1986.

[MPLW22] Jieming Mao, Renato Paes Leme, and Kangning Wang. Interactive communication in bilateral trade. In *13th Innovations in Theoretical Computer Science Conference, ITCS*, pages 105:1–105:21, 2022.

[MS83] Roger B. Myerson and Mark A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of economic theory*, 29(2):265–281, 1983.

[MSW21] Kamesh Munagala, Zeyu Shen, and Kangning Wang. Optimal algorithms for multiwinner elections and the chamberlin-courant rule. In *EC*, pages 697–717, 2021.

[MSWW22] Kamesh Munagala, Yiheng Shen, Kangning Wang, and Zhiyi Wang. Approximate core for committee selection via multilinear extension and market clearing. In *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9 - 12, 2022*, pages 2229–2252, 2022.

[MU05] Michael Mitzenmacher and Eli Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis.* Cambridge University Press, 2005.

[MW19] Kamesh Munagala and Kangning Wang. Improved metric distortion for deterministic social choice rules. In *Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019, Phoenix, AZ, USA, June 24-28, 2019*, pages 245–262, 2019.

[Mye81] Roger B Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

[Nas50] John F. Nash. The bargaining problem. *Econometrica*, 18(2):155–162, 1950.

[NWF78] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An analysis of approximations for maximizing submodular set functions—I. *Mathematical Programming*, 14(1):265–294, 1978.

[pbs] The Stanford Participatory Budgeting Platform.

[PPS21] Dominik Peters, Grzegorz Pierczyński, and Piotr Skowron. Proportional participatory budgeting with additive utilities. In *NeurIPS*, 2021.

[PR06] Ariel D Procaccia and Jeffrey S Rosenschein. The distortion of cardinal preferences in voting. In *International Workshop on Cooperative Information Agents*, pages 317–331. Springer, 2006.

[PS97] Alessandro Panconesi and Aravind Srinivasan. Randomized distributed edge coloring via an extension of the chernoff-hoeffding bounds. *SIAM J. Comput.*, 26(2):350–368, 1997.

[PS20] Dominik Peters and Piotr Skowron. Proportionality and the limits of welfarism. In *EC*, pages 793–794, 2020.

[Sat75] Mark Allen Satterthwaite. Strategy-proofness and arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of economic theory*, 10(2):187–217, 1975.

[Sch07] Norman Schofield. *The spatial model of politics*. Routledge, 2007.

[Sch11] Markus Schulze. A new monotonic, clone-independent, reversal symmetric, and condorcet-consistent single-winner election method. *Social Choice and Welfare*, 36(2):267–303, 2011.

[SE17] Piotr Krzysztof Skowron and Edith Elkind. Social choice under metric preferences: Scoring rules and STV. In *Proceedings of the 31th AAAI Conference on Artificial Intelligence (AAAI)*, pages 706–712, 2017.

[SFL16] Piotr Skowron, Piotr Faliszewski, and Jerome Lang. Finding a collective set of items: From proportional multirepresentation to group recommendation. *Artificial Intelligence*, 241:191 – 216, 2016.

[SFS15] Piotr Skowron, Piotr Faliszewski, and Arkadii Slinko. Achieving fully proportional representation: Approximability results. *Artificial Intelligence*, 222:67 – 103, 2015.

[STZ18] Weiran Shen, Pingzhong Tang, and Yulong Zeng. A closed-form characterization of buyer signaling schemes in monopoly pricing. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS*, pages 1531–1539, 2018.

[STZ19] Weiran Shen, Pingzhong Tang, and Yulong Zeng. Buyer signaling games in auctions. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS*, pages 1591–1599, 2019.

[Thi95] Thorvald N Thiele. Om flerfoldsvalg. *Oversigt over det Kongelige Danske Videnskabernes Selskabs Forhandlinger*, 1895:415–441, 1895.

[Tid87] T Nicolaus Tideman. Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4(3):185–206, 1987.

[TR00] Nicolaus Tideman and Daniel Richardson. Better voting methods through technology: The refinement-manageability trade-off in the single transferable vote. *Public Choice*, 103(1):13–34, 2000.

[Von08] Jan Vondrák. Optimal approximation for the submodular welfare problem in the value oracle model. In *STOC*, page 67–74, 2008.

[XRDT15] Haifeng Xu, Zinovi Rabinovich, Shaddin Dughmi, and Milind Tambe. Exploring information asymmetry in two-stage security games. In *29th AAAI Conference on Artificial Intelligence, AAAI*, 2015.

[Yan11] Qiqi Yan. Mechanism design via correlation gap. In *SODA*, page 710–719, 2011.

Biography

Kangning Wang is a Ph.D. candidate in Computer Science at Duke University. He is advised by Kamesh Munagala. He received his bachelor's degree from Institute for Interdisciplinary Information Sciences (Yao Class), Tsinghua University in 2017. He received Best Paper Award at WINE 2018, and Best Preliminary Exam Award from Duke Computer Science in 2020.