

Optimized Distortion and Proportional Fairness in Voting

SOROUGH EBADIAN, University of Toronto

ANSON KAHNG, University of Toronto

DOMINIK PETERS, University of Toronto and CNRS, Université Paris Dauphine–PSL

NISARG SHAH, University of Toronto

A voting rule decides on a probability distribution over a set of m alternatives, based on rankings of those alternatives provided by agents. We assume that agents have cardinal utility functions over the alternatives, but voting rules have access to only the rankings induced by these utilities. We evaluate how well voting rules do on measures of social welfare and of proportional fairness, computed based on the hidden utility functions.

In particular, we study the *distortion* of voting rules, which is a worst-case measure. It is an approximation ratio comparing the utilitarian social welfare of the optimum outcome to the social welfare produced by the outcome selected by the voting rule, in the worst case over possible input profiles and utility functions that are consistent with the input. The previous literature has studied distortion with unit-sum utility functions (which are normalized to sum to 1), and left a small asymptotic gap in the best possible distortion. Using tools from the theory of fair multi-winner elections, we propose the first voting rule which achieves the optimal distortion $\Theta(\sqrt{m})$ for unit-sum utilities. Our voting rule also achieves optimum $\Theta(\sqrt{m})$ distortion for a larger class of utilities, including unit-range and approval (0/1) utilities.

We then take a similar worst-case approach to a quantitative measure of the fairness of a voting rule, called *proportional fairness*. Informally, it measures whether the influence of cohesive groups of agents on the voting outcome is proportional to the group size. We show that there is a voting rule which, without knowledge of the utilities, can achieve an $O(\log m)$ -approximation to proportional fairness, which is the best possible approximation. As a consequence of its proportional fairness, we show that this voting rule achieves $O(\log m)$ distortion with respect to the Nash welfare, and selects a distribution that is approximately stable by being an $O(\log m)$ -approximation to the core, making it interesting for applications in participatory budgeting.

CCS Concepts: • **Theory of computation** → **Design and analysis of algorithms**; *Approximation algorithms analysis*; • **Applied computing** → *Economics*; *Multi-criterion optimization and decision-making*.

Additional Key Words and Phrases: computational social choice, voting, participatory budgeting, distortion, Nash welfare, social welfare, core, proportional fairness

ACM Reference Format:

Soroush Ebadian, Anson Kahng, Dominik Peters, and Nisarg Shah. 2022. Optimized Distortion and Proportional Fairness in Voting. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC '22)*, July 11–15, 2022, Boulder, CO, USA. ACM, New York, NY, USA, 38 pages. <https://doi.org/10.1145/3490486.3538339>

1 INTRODUCTION

We consider the problem of designing voting rules that aggregate agents' ranked preferences and arrive at a collective decision with high social welfare and which is fair to all agents. Throughout, we focus on probabilistic voting rules, which take as input a preference profile of complete rankings of a set A of m alternatives and output a probability distribution over A .

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC '22, July 11–15, 2022, Boulder, CO, USA.

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9150-4/22/07...\$15.00

<https://doi.org/10.1145/3490486.3538339>

	Distortion, unit-sum	Distortion, unit-range	Proportional fairness
Stable Lottery Rule	$\Theta(\sqrt{m})$	$\Theta(\sqrt{m})$	—
Stable Committee Rule	$\Theta(\sqrt{m})$	$\Theta(\sqrt{m})$	$O(\sqrt{m})$
Harmonic Rule (truthful)	$\Theta(\sqrt{m \log m})$	$\Theta(m^{2/3} \log^{1/3} m)$	$\Theta(\sqrt{m \log m})$
Best possible	$\Theta(\sqrt{m})$	$\Theta(\sqrt{m})$	$\Theta(\log m)$
Best possible, truthful	$\Theta(\sqrt{m \log m})$	$\Theta(m^{2/3})$	$\Omega(\sqrt{m}), O(\sqrt{m \log m})$

Table 1. Overview of our results. Previously known results: $O(\sqrt{m \log m})$ upper bound on the distortion of the harmonic rule for unit-sum utilities [Boutilier et al. 2015], $\Omega(\sqrt{m})$ [Boutilier et al. 2015] and $\Omega(\sqrt{m \log m})$ [Bhaskar et al. 2018] lower bounds on the distortion of non-truthful and truthful rules for unit-sum utilities, respectively, and $\Theta(m^{2/3})$ distortion of the best possible truthful rule for unit-range utilities [Filos-Ratsikas and Miltersen 2014; Lee 2013].

In order to evaluate the social welfare and fairness of voting rules, we build upon the framework of *implicit utilitarian voting* [Procaccia and Rosenschein 2006], which assumes that each agent i has a cardinal utility function $u_i : A \rightarrow \mathbb{R}_{\geq 0}$ over alternatives, but reports only the induced ranking over alternatives to the voting rule. While in principle a voting rule could elicit the precise utility values, it is more common in the literature to ask for rankings. This makes for a simple elicitation protocol, which can ease the cognitive burden on agents (because they need not precisely determine their own utility values), and preserves the privacy of any agents who may not wish to reveal their exact utilities to a voting rule.

The implicit utilitarian framework allows us to quantify the efficiency of a given voting rule: Given an input profile of rankings, we can measure efficiency as the worst-case ratio between the social welfare of the optimal outcome and the social welfare of the outcome selected by the voting rule, where the worst case is taken over all utility functions consistent with the ordinal rankings reported to the voting rule. This quantity is known as *distortion* and has been widely studied. The existing literature commonly analyzes distortion for the class of *unit-sum utilities*, in which each agent’s total utility is normalized to 1 [Boutilier et al. 2015; Caragiannis et al. 2017; Filos-Ratsikas et al. 2020; Mandal et al. 2019, 2020]. For deterministic voting rules, which place the entire probability mass of 1 on a single alternative, it is known that the best possible distortion is $\Theta(m^2)$ [Caragiannis et al. 2017; Caragiannis and Procaccia 2011]. For probabilistic voting rules, which we are interested in, Boutilier et al. [2015] prove a lower bound of $\Omega(\sqrt{m})$. We propose the first voting rule achieving the asymptotically optimal distortion of $O(\sqrt{m})$, matching their lower bound and resolving an important open question in this line of work. Our proof shows that the same rule is also optimal for *unit-range* utilities (which are normalized to range between 0 and 1) with the same $O(\sqrt{m})$ distortion. This improves upon the previous best-known distortion of $O(m^{2/3})$ [Filos-Ratsikas 2015; Lee 2013]. This $O(\sqrt{m})$ distortion of our rule is also optimal for the special case of *approval utilities*, in which each agent has utility 1 for a subset of alternatives and utility 0 for the rest. This class corresponds naturally to approval voting but, to the best of our knowledge, has not been studied in the context of distortion.¹ Our rule can be computed in polynomial time.

¹Distortion for approval utilities makes sense in contexts where agents may find it easier to rank alternatives than to assign them approval utilities. For example, if the alternatives are budget divisions, a project leader would naturally rank the divisions by the amount of money allocated to their project. But the eventual utility depends on whether the money is enough to deliver the project or not, and the required amount may be unknown at the time of voting.

Interestingly, while our new voting rule achieves low distortion (i.e., high social welfare), it internally aims for a fair outcome. Specifically, it uses tools from multi-winner voting for selecting a committee (a fixed-size subset of alternatives) that is *representative*. Informally, as many agents as possible should have one of their highly-ranked alternatives in the committee. There is an intuitive case for considering representative committees for achieving low distortion: Suppose a voting rule places little weight on the highly-ranked alternatives of a large group of agents. Then the voting rule may incur high distortion when those agents feel strongly about their preferences and all other agents are indifferent. This suggests that, at least in some settings, if you want to be efficient, it pays to also be fair.

While we use fair committees as a means to achieve high social welfare, we are also interested in fairness as an end. We wish to achieve a notion of fairness defined for our single-winner setting. Specifically, we adapt a quantitative measure from network theory called *proportional fairness* to the voting context. This measure is phrased in terms of agents’ utility functions, and so we combine it with the worst-case philosophy of distortion to obtain a way to measure the fairness of voting rules. Intuitively, for an outcome to do well with respect to proportional fairness, it cannot be the case that any large group of agents gets too little utility, where “too little” is a function of how large the group is and how easy it is to give high utility to the group.

If we knew the underlying agent utilities, we could compute a distribution that is 1-proportionally fair. We show that given only a preference profile of rankings, there always exists a distribution that is $O(\log m)$ -proportionally fair regardless of the underlying utility functions (consistent with the input rankings). Our existence proof uses the minimax theorem for zero-sum games. We show that our result is optimal because there are preference profiles for which every distribution has an approximation to proportional fairness that is no better than $\Omega(\log m)$ under some consistent utility functions. We then show that, given a preference profile, we can compute in polynomial time a distribution with an (almost) optimal approximation to proportional fairness.

Proportional fairness is an interesting measure because voting rules that do well on it automatically do well on other fairness measures as well. For example, it is widely recognized that maximizing the *Nash welfare* instead of the utilitarian welfare gives fairer outcomes (the Nash welfare of an outcome is the *product* of agent utilities instead of the sum). We can define a version of distortion for the Nash welfare, and our rule for proportional fairness will guarantee $O(\log m)$ distortion for this objective.² Another fairness property is taken from the literature on *participatory budgeting* (PB) [Fain et al. 2016]. We can interpret a probabilistic voting rule as dividing a fixed budget between different projects, and agents vote by ranking those projects. Agents wish to see more money spent on the projects they rank higher. An important goal in PB is to provide proportional representation in that any $x\%$ of the population cannot find an allocation of $x\%$ of the budget which provides them a Pareto improvement (i.e., does not hurt any of them and strictly improves some). This aim can be formalized using the concept of the core. Our rule for proportional fairness selects an outcome that provides an $O(\log m)$ -approximation to the core.

Table 1 provides an up-to-date account of the results on distortion for unit-sum and unit-range utilities as well as proportional fairness.

1.1 Related Work

There are many papers that study the distortion of voting rules, beginning with the work of Procaccia and Rosenschein [2006], who analyze the distortion of many common voting rules. Caragiannis and Procaccia [2011] also evaluate the distortion of prominent voting rules, but from the perspective of optimizing embeddings, which (perhaps randomly) map cardinal utilities to ordinal votes that

²We show a constant lower bound of e for the distortion for Nash welfare, thus leaving a substantial gap.

voting rules take as input. Their work, together with that of Caragiannis et al. [2017], identifies the best possible distortion via deterministic voting rules to be $\Theta(m^2)$. Boutilier et al. [2015] study probabilistic voting rules and derive a lower bound of $\Omega(\sqrt{m})$ on the optimal distortion for probabilistic voting rules with unit-sum utilities. They also design an artificial rule (tailored specifically to the unit-sum distortion context) which establishes an upper bound of $O(\sqrt{m} \log^* m)$.³ Our $O(\sqrt{m})$ upper bound matches their lower bound and eliminates the $\log^* m$ gap.

Boutilier et al. [2015] also propose the *harmonic rule* based on the harmonic scoring rule and show that it achieves $O(\sqrt{m \log m})$ distortion for unit-sum utilities. Bhaskar et al. [2018] point out that this voting rule is truthful (in expectation with respect to any consistent utility function), and prove that any truthful rule must incur at least $\Omega(\sqrt{m \log m})$ distortion, making the harmonic rule asymptotically optimal, subject to truthfulness. Distortion subject to truthfulness had first been studied by Filos-Ratsikas and Miltersen [2014], who consider unit-range utilities and prove that any truthful rule must incur at least $\Omega(m^{2/3})$ distortion. Their proof also implies this bound for approval utilities. Lee [2013] proposed a truthful method that achieves a matching upper bound for unit-range utilities (see also the work of Filos-Ratsikas [2015]). In the appendix, we show that the harmonic rule achieves distortion $\Theta(m^{2/3} \log^{1/3} m)$ for approval and for unit-range utilities, matching the lower bound subject to truthfulness up to a logarithmic factor. Using the techniques of Filos-Ratsikas and Miltersen [2014] and Bhaskar et al. [2018], we derive a lower bound of $\Omega(\sqrt{m})$ on our proportional fairness objective subject to truthfulness, and show that the harmonic rule again matches this, up to a logarithmic factor.

Implicit utilitarian voting can be seen as a protocol for reducing communication complexity by asking agents to report ordinal preferences in place of cardinal utilities, so it is natural to study the trade-off between the communication complexity (the number of bits of information each agent is asked to report) and the optimal distortion achievable. Mandal et al. [2019, 2020] characterize the Pareto frontier of this tradeoff, showing that in order to achieve distortion d , probabilistic voting rules require agents to communicate only $\tilde{\Theta}(m/d^3)$ bits of information whereas deterministic voting rules require $\Theta(m/d)$ bits, establishing probabilistic rules as superior in this context. Amanatidis et al. [2021] considered making a few value queries (asking agents to report their utility for an alternative) or comparison queries (asking agents to report whether the ratio of their utilities for two alternatives is at least a threshold) on top of their reported ordinal preferences. They prove that asking only $O(\log^2 m)$ value queries or $O(\log^2 m)$ comparison queries is sufficient to achieve constant distortion.

Going beyond single-winner voting, Caragiannis et al. [2017] study distortion (and another closely related objective called regret) for multi-winner voting, where the goal is to select a committee of k alternatives for a given size k . They assume that the utility of an agent for a committee is the maximum utility of the agent for any alternative in the committee. They prove that the optimal distortion of deterministic rules is $\Theta(1 + m(m - k)/k)$, implying an optimal distortion of $\Theta(m^2)$ for deterministic single-winner voting. For probabilistic rules, they leave a gap of $\Theta(m^{1/6})$ between their upper and lower bounds for the optimal distortion. Recently, Borodin et al. [2022] close this gap by building upon our work. They extend our single-winner rule with $O(\sqrt{m})$ distortion to multi-winner voting and prove that it achieves the optimal distortion of $\Theta(\min(\sqrt{m}, m/k))$.

Benade et al. [2021] study participatory budgeting, which is an extension of multi-winner voting in which each alternative has a cost and the goal is to find a subset of alternatives with total cost at most a given budget. They focus on a different utility model, in which the utility of an agent for a set of alternatives is the sum of her utilities for the alternatives in the set. They compare four protocols for eliciting agent preferences and prove that while ranked preferences lead to $O(\sqrt{m} \cdot \log m)$

³ $\log^* m$ is the *iterated logarithm* of m , which is the number of times \log needs to be applied to m until the result is at most 1.

distortion with probabilistic aggregation, threshold approval votes, which ask agents to identify alternatives for which their utility is at least a specified threshold, lead to a significantly lower distortion of $O(\log^2 m)$. Bhaskar et al. [2018] show that the near-optimal $O(\sqrt{m} \cdot \log m)$ distortion for participatory budgeting with ranked preferences can in fact be obtained via a truthful voting rule, establishing that truthfulness comes at minimal cost even in this general model.

In all these papers, agents are modeled to have *normalized utilities* for alternatives. Initiated by the work of Anshelevich et al. [2018], a large recent literature about *metric distortion* instead models agents as having *costs* for alternatives. This literature makes the assumption that the cost of an agent for an alternative is the distance between them in an underlying metric space, and aims to approximate the utilitarian social cost (i.e., the sum of agent costs) [Anshelevich et al. 2018; Anshelevich and Postl 2017; Anshelevich and Sekar 2016; Caragiannis et al. 2022; Munagala and Wang 2019]. It turns out that the metric structure allows significantly better distortion bounds: the best distortion of deterministic rules is 3 [Gkatzelis et al. 2020] (compared to $\Theta(m^2)$ in the non-metric setting) and that for probabilistic rules is between 2.0261 and $3 - 2/m$ [Charikar and Ramakrishnan 2022; Kempe 2020] (compared to $\Theta(\sqrt{m})$ distortion in the non-metric setting). Note that probabilistic rules are superior to deterministic rules in the metric setting as well.

Fairness of single-winner voting rules has received less attention than distortion. For probabilistic voting rules, fairness has been studied in a series of papers that interpret the output distribution as a division of a budget. Most work has studied this in a model with known approval utilities of the agents [Aziz et al. 2019; Bogomolnaia et al. 2005; Brandl et al. 2021; Duddy 2015]. Airiau et al. [2019] study probabilistic voting rules which take as input ranked preferences, and then convert those preferences into utilities using a fixed scoring vector (such as Borda). The rules then maximize the Nash welfare (the geometric mean of agent utilities) or the egalitarian welfare (the minimum agent utility) and its leximin refinement. Note that in our work the utilities are unknown. They prove that Nash-welfare-based rules satisfy the *SD-core*. This is a weaker axiom than the core that we introduce in Section 2, which in the terminology of Aziz et al. [2018] could be called the *strong SD-core*. We note that SD-core implies no better than an m -approximation of our (strong) core (for example, random dictatorship satisfies SD-core and is in the m -approximate core), whereas we achieve an $O(\log m)$ -approximate core. In a model where voters report their utilities, Fain et al. [2016] investigate the core and propose a polynomial-time algorithm for finding an outcome in the core via the so-called Lindahl equilibrium. Note that they do not require an approximation to the core because utilities are known. They also point out connections to proportional fairness.

Fairness in voting has been studied in detail for deterministic multi-winner voting rules. Various fairness notions have been studied that require every group of agents to have representation in the committee, with larger and more cohesive groups having better representation. This includes notions such as justified representation (JR), extended justified representation (EJR) [Aziz et al. 2017a], proportional justified representation (PJR) [Sánchez-Fernández et al. 2017], full justified representation (FJR) [Peters et al. 2021a] and the proportionality degree [Skowron 2021]. Cheng et al. [2020] prove that there always exists a distribution over committees that satisfies a stronger fairness notion called *stability*; this is the main tool we use to achieve $O(\sqrt{m})$ distortion for single-winner voting. Jiang et al. [2020] derandomize this result to prove that there always exists a committee satisfying 32-approximate stability; we show that this derandomized result can be used to achieve $O(\sqrt{m})$ distortion with respect to the Nash welfare, but we are able to improve on that bound to achieve $O(\log m)$ distortion using the minimax theorem. Fain et al. [2018] study a more general model of public goods and achieve different approximations to the core under various constraints on feasible outcomes.

2 PRELIMINARIES

For $t \in \mathbb{N}$, let $[t] = \{1, \dots, t\}$. For a set X , let $\Delta(X)$ be the set of probability distributions \mathbf{x} over X .

Voting. Let N be a set of n agents and A be a set of m alternatives. For $k \in [m]$, let $\mathcal{P}_k(A)$ denote the set of all subsets of A of size k . Each agent $i \in N$ submits a *preference ranking* over the alternatives, encoded by a bijective rank function $\sigma_i : A \rightarrow [m]$. For example, if $\sigma_i(a) = 1$, then a is the most-preferred alternative for agent i . We use $a \succ_i a'$ to denote $\sigma_i(a) < \sigma_i(a')$ (agent i ranks a strictly above a') and $a \succsim_i a'$ to denote $\sigma_i(a) \leq \sigma_i(a')$. We refer to the collection $\vec{\sigma} = (\sigma_i)_{i \in N}$ as a *preference profile*. A (probabilistic) *voting rule* f is a function that takes a preference profile $\vec{\sigma}$ as input and outputs a distribution over alternatives. Note that the output of a voting rule can be interpreted as a randomized selection of alternatives, but also as a division of some divisible resource (such as time or a budget) between the alternatives.

Utilities. A *utility function* $u : A \rightarrow \mathbb{R}_{\geq 0}$ assigns a non-negative utility to each alternative. We can extend u to also assign utility values to distributions $\mathbf{x} \in \Delta(A)$ over alternatives by setting $u(\mathbf{x}) = \mathbb{E}_{a \sim \mathbf{x}} u(a)$. We assume that when agents submit ranked preferences, they have more expressive underlying cardinal preferences. Given a preference profile $\vec{\sigma}$, we say that a utility function u_i for agent i is *consistent* with her preference ranking if for all $a, a' \in A$ such that $a \succ_i a'$, we have $u_i(a) \geq u_i(a')$. Note that we allow alternatives to have equal utility, and then the agent can break ties arbitrarily when submitting a preference ranking. We refer to a collection $\vec{u} = (u_i)_{i \in N}$ as a *utility profile*. We use the notation $\vec{u} \triangleright \vec{\sigma}$ to indicate that u_i is consistent with σ_i for each agent i . Note that voting rules have access to the preference profile but not to the utility profile.

Utility classes. Let \mathcal{U}^{all} denote the class of all possible utility functions. We also study the following standard restricted utility classes.

- $\mathcal{U}^{\text{unit-sum}}$ is the class of *unit-sum* utility functions u satisfying $\sum_{a \in A} u(a) = 1$.
- $\mathcal{U}^{\text{unit-range}}$ is the class of *unit-range* utility functions u satisfying $\max_{a \in A} u(a) = 1$.⁴
- $\mathcal{U}^{\text{approval}}$ is the class of *approval* utility functions u satisfying $u(a) \in \{0, 1\}$ for all $a \in A$ and $u(a) = 1$ for at least one $a \in A$.

We introduce a new class of *balanced* utility functions, where the highest utility intensity that can be expressed is at most 1, and where the total utility of the utility function is at least 1.

- $\mathcal{U}^{\text{balanced}}$ is the class of utility functions u satisfying $u(a) \leq 1$ for all $a \in A$ and $\sum_{a \in A} u(a) \geq 1$.

Note that $\mathcal{U}^{\text{unit-sum}} \subseteq \mathcal{U}^{\text{balanced}}$ and $\mathcal{U}^{\text{approval}} \subseteq \mathcal{U}^{\text{unit-range}} \subseteq \mathcal{U}^{\text{balanced}}$. Our positive distortion result will work for the entire class of balanced utility functions.

In this work, we focus on two metrics for evaluating voting rules: distortion, which is a measure of social welfare, and proportional fairness, which is a measure of fairness.

2.1 Distortion

Given the utility profile \vec{u} , the *utilitarian welfare* of a distribution over alternatives $\mathbf{x} \in \Delta(A)$ is defined as $\text{UW}(\mathbf{x}, \vec{u}) = \sum_{i \in N} u_i(\mathbf{x})$.

If one could observe the underlying utilities, an argument dating back to Bentham [1789] suggests that picking the alternative maximizing the utilitarian welfare is the best choice. However, a voting rule is allowed to observe only the preference profile $\vec{\sigma}$, thus obtaining partial information about the utility profile \vec{u} . In this case, we measure the efficiency of the voting rule by the worst-case approximation ratio it achieves for maximizing the utilitarian welfare.

⁴Some definitions of unit-range utilities require $\min_{a \in A} u(a) = 0$ in addition, but this is not necessary for our results.

		Agent i_1	Agent i_2	Agent i_3
Preferences	$\vec{\sigma}$	$a_1 > a_2 > a_3$	$a_2 > a_1 > a_3$	$a_1 > a_3 > a_2$
Utilities	\vec{u}_1	$1/2, 1/3, 1/6$	$1/2, 1/2, 0$	$1/3, 1/3, 1/3$
	\vec{u}_2	$1/2, 1/2, 0$	$1, 0, 0$	$1/3, 1/3, 1/3$

Table 2. A preference profile with three agents and three alternatives along with two consistent unit-sum utility profiles.

Definition 2.1 (Distortion). Given a utility profile \vec{u} , the *distortion* of a distribution $\mathbf{x} \in \Delta(A)$ is the ratio between the highest possible social welfare and the social welfare of \mathbf{x} under \vec{u} :

$$D(\mathbf{x}, \vec{u}) = \frac{\max_{\mathbf{y} \in \Delta(A)} \text{UW}(\mathbf{y}, \vec{u})}{\text{UW}(\mathbf{x}, \vec{u})}.$$

The distortion of \mathbf{x} on a preference profile $\vec{\sigma}$ for a utility class \mathcal{U} is obtained by taking the worst case over all utility profiles $\vec{u} \in \mathcal{U}^n$ consistent with $\vec{\sigma}$.

$$D(\mathbf{x}, \vec{\sigma}, \mathcal{U}) = \sup_{\vec{u} \in \mathcal{U}^n: \vec{u} \triangleright \vec{\sigma}} D(\mathbf{x}, \vec{u}).$$

Given a number m of alternatives, the distortion of a voting rule f for utility class \mathcal{U} is $D_m(f, \mathcal{U}) = \sup_{\vec{\sigma}} D(f(\vec{\sigma}), \vec{\sigma}, \mathcal{U})$, where the supremum is taken over all preference profiles $\vec{\sigma}$ with m alternatives and any number of agents.

Example 2.2. Table 2 shows a preference profile with three agents and three alternatives. Consider the distribution $\mathbf{x} = (a_1 : 1/2, a_2 : 1/4, a_3 : 1/4)$. Let us evaluate its distortion under two utility profiles given in the table.

- For utility profile \vec{u}_1 , the social welfare of \mathbf{x} is $\text{UW}(\mathbf{x}, \vec{u}_1) = 3/8 + 3/8 + 1/3 = 13/12$. For \vec{u}_1 , the optimal outcome is $\mathbf{y} = (a_1 : 1)$ with $\text{UW}(\mathbf{y}, \vec{u}_1) = 4/3$. Hence, $D(\mathbf{x}, \vec{u}_1) = \frac{4/3}{13/12} \approx 1.23$.
- For utility profile \vec{u}_2 , the social welfare of \mathbf{x} is $\text{UW}(\mathbf{x}, \vec{u}_2) = 5/12 + 11/24 + 1/12 = 23/24$. For \vec{u}_2 , the optimal outcome is $\mathbf{y} = (a_1 : 1)$ with $\text{UW}(\mathbf{y}, \vec{u}_2) = 11/6$. Hence, $D(\mathbf{x}, \vec{u}_2) = \frac{11/6}{23/24} = 44/23 \approx 1.91$.

Thus, under \vec{u}_1 , it is possible to obtain 23% more social welfare than \mathbf{x} , and under \vec{u}_2 , it is possible to obtain 91% more. Using a simple linear program, one can check that \vec{u}_2 is the worst case for utility profiles from $\mathcal{U}^{\text{unit-sum}}$, so $D(\mathbf{x}, \vec{\sigma}, \mathcal{U}^{\text{unit-sum}}) = 44/23 \approx 1.91$. Using a more sophisticated linear program [Boutlier et al. 2015], one can find the distribution with the lowest possible distortion for unit-sum utilities, which on this profile is $\mathbf{x}^* \approx (a_1 : 0.5882, a_2 : 0.4118, a_3 : 0)$, achieving distortion of approximately 1.54.

As we mention in Example 2.2, given a preference profile $\vec{\sigma}$, one can find a distribution \mathbf{x} minimizing $D(\mathbf{x}, \vec{\sigma}, \mathcal{U}^{\text{unit-sum}})$ by solving a linear program proposed by Boutlier et al. [2015]. Their approach works for any utility class that is described by linear constraints, so it can be used to find instance-optimal distributions for unit-sum, for unit-range, and for balanced utilities. We show in Lemma B.1 in the appendix that the distribution minimizing distortion for unit-range utilities also minimizes distortion for approval utilities, and so instance-optimal distributions for approval utilities can also be found in polynomial time.

We have defined the distortion for a class of utility functions \mathcal{U} by taking the worst case over all utility profiles \vec{u} in which the utility function u_i of every agent i belongs to \mathcal{U} . Most naturally, one would like to analyze the distortion for the class of all utility functions \mathcal{U}^{all} . However, the worst-case distortion for this class is degenerate: the rule that always selects the uniform distribution has $O(m)$ distortion, and it is easy to see that any rule has at least $\Omega(m)$ distortion (by considering

utility profiles where some agents care a lot and others not at all). Thus, without some additional restrictions on cardinal utilities (such as unit-sum or unit-range), it turns out that ordinal preferences do not provide significant information about utilitarian welfare.

2.2 Nash Welfare Distortion

Distortion is typically defined with respect to utilitarian welfare, but the same principle can be applied to other welfare functions. We will in particular study *Nash welfare* (NW), which is the geometric mean of agent utilities: $NW(\mathbf{x}, \vec{u}) = (\prod_{i \in N} u_i(\mathbf{x}))^{1/n}$. We can define the distortion $D_m^{\text{NW}}(f, \mathcal{U})$ of a voting rule f for Nash welfare by replacing the utilitarian welfare UW in Definition 2.1 by NW.

Nash welfare is sometimes viewed as a combined measure of efficiency and fairness. It measures efficiency in a Pareto sense (if everyone’s utility increases then so does Nash welfare), and it measures fairness because if some agent has very low utility then this has a strong negative impact on overall Nash welfare.

Remarkably, the Nash welfare is *scale invariant*, i.e., multiplying the utility function of an agent by some factor does not change the comparison between the Nash welfare of two distributions over alternatives. Hence, we have that $D_m^{\text{NW}}(f, \mathcal{U}^{\text{all}}) = D_m^{\text{NW}}(f, \mathcal{U}^{\text{unit-sum}}) = D_m^{\text{NW}}(f, \mathcal{U}^{\text{unit-range}})$ for every voting rule f .

2.3 Core

When we view a distribution \mathbf{x} as a division of a budget between the alternatives, the core is a fairness axiom that intuitively guarantees every group of agents an influence proportional to its size, provided the agents in the group have similar enough preferences.

Let $\alpha \geq 1$. We will define an α -approximate notion of the core which coincides with the standard version when $\alpha = 1$. Similar α -approximations to the core have been studied in discrete budgeting settings [Fain et al. 2018; Munagala et al. 2022; Peters and Skowron 2020]. A distribution over alternatives $\mathbf{x} \in \Delta(A)$ is said to be in the α -core with respect to utility profile \vec{u} if there is no subset of agents S and distribution over alternatives $\mathbf{y} \in \Delta(A)$ such that

$$\frac{|S|}{|N|} \cdot u_i(\mathbf{y}) \geq \alpha \cdot u_i(\mathbf{x}) \text{ for every agent } i \in S,$$

and at least one of these inequalities is strict.⁵ A voting rule f is said to be in the α -core if, for every preference profile $\vec{\sigma}$, $f(\vec{\sigma})$ is in the α -core with respect to every utility profile \vec{u} consistent with $\vec{\sigma}$.

The core notion is inspired by cooperative game theory, where it is seen as a stability notion: if a group of agents is not treated fairly, then those agents can leave the system in order to use their fraction of the budget in a preferable way. The core is typically defined in settings where agents’ utility functions are known; our notion is phrased for the case where the rule has access to only ordinal information. Note that to be in the α -core, a rule needs to avoid deviations for all consistent utilities. This makes sense as a conservative stability notion because agents will presumably make their decision to leave based on their actual utilities.

While 1-core can be achieved when exact utilities are known (see Section 2.4), it is easy to see that no rule can satisfy 1-core given only the ordinal preferences. For example, consider the profile from Example 2.2 with preferences $(a_1 > a_2 > a_3, a_2 > a_1 > a_3, a_1 > a_3 > a_2)$. For the utility profile where each agent approves just their top alternative, there is a unique distribution \mathbf{x} that satisfies the 1-core, namely $\mathbf{x} = (a_1 : 2/3, a_2 : 1/3, a_3 : 0)$. This is because if $\mathbf{x}(a_1)$ was any lower, then the first and third agents could deviate; if $\mathbf{x}(a_2)$ was lower, then the second agent could deviate. However, \mathbf{x} fails the 1-core if we change the utility profile so that the second agent gives the same utility to all

⁵Equivalently, this condition requires that there is no set $S \subseteq N$ and *partial* distribution $\mathbf{y} : A \rightarrow [0, 1]$ with $\sum_a \mathbf{y}_a = |S|/n$ such that $u_i(\mathbf{y}) \geq \alpha \cdot u_i(\mathbf{x})$ for every agent $i \in S$ and at least one of these inequalities is strict [Fain et al. 2016, 2018].

alternatives. For that utility profile, all three agents can deviate together by proposing to place the entire budget on a_1 .

2.4 Proportional Fairness

We have now seen two notions that are connected to fairness (Nash welfare and the core). A third such notion is *proportional fairness*, which was first proposed in communication networks [Kelly et al. 1998] but is easily adapted to social choice more generally. This is a quantitative way of measuring the fairness of a distribution. As we will see, proportional fairness is intimately connected to the other two notions.

Definition 2.3 (Proportional Fairness). Let $\mathbf{x} \in \Delta(A)$ be a distribution over alternatives. Given a utility profile \vec{u} , we write

$$\text{PF}(\mathbf{x}, \vec{u}) = \max_{\mathbf{y} \in \Delta(A)} \frac{1}{n} \sum_{i \in N} \frac{u_i(\mathbf{y})}{u_i(\mathbf{x})} = \max_{a \in A} \frac{1}{n} \sum_{i \in N} \frac{u_i(a)}{u_i(\mathbf{x})}. \quad (1)$$

Given a preference profile $\vec{\sigma}$ and a utility class \mathcal{U} , we write

$$\text{PF}(\mathbf{x}, \vec{\sigma}, \mathcal{U}) = \sup_{\vec{u} \in \mathcal{U}^n: \vec{u} \succ \vec{\sigma}} \text{PF}(\mathbf{x}, \vec{u}).$$

If $\text{PF}(\mathbf{x}, \vec{\sigma}, \mathcal{U}) \leq \alpha$, we say that \mathbf{x} is α -proportionally fair (for utility class \mathcal{U}).

Given a number m of alternatives, we say that a voting rule f is α -proportionally fair (for utility class \mathcal{U}) if for every preference profile $\vec{\sigma}$ with m alternatives and any number of agents, we have $\text{PF}(f(\vec{\sigma}), \vec{\sigma}, \mathcal{U}) \leq \alpha$. We also write $\text{PF}_m(f, \mathcal{U}) = \sup_{\vec{\sigma}} \text{PF}(f(\vec{\sigma}), \vec{\sigma}, \mathcal{U})$ so that f is α -proportionally fair if and only if $\text{PF}_m(f, \mathcal{U}) \leq \alpha$.

Example 2.4. Take the example in Table 2. For distribution $\mathbf{x} = (a_1 : 1/2, a_2 : 1/4, a_3 : 1/4)$ and utilities \vec{u}_1 , we have $\text{PF}(\mathbf{x}, \vec{u}_1) = \frac{1}{3} \max\{11/3, 29/9, 13/9\} = 11/9 \approx 1.22$. For utilities \vec{u}_2 , we have $\text{PF}(\mathbf{x}, \vec{u}_2) = \frac{1}{3} \max\{7/3, 19/3, 1\} = 19/9 \approx 2.11$. Using Lemma 4.1 in Section 4, we can check that \vec{u}_2 is the worst-case utility profile for distribution \mathbf{x} , so that $\text{PF}(\mathbf{x}, \vec{u}) \leq 19/9 \approx 2.11$ for all utility profiles $\vec{u} \in \mathcal{U}^{\text{all}}$. Hence \mathbf{x} is 2.11-proportionally fair.

Using the techniques described in Section 4.3, we can establish that the optimal distribution with respect to proportional fairness is $\mathbf{x}^* \approx (a_1 : 0.586, a_2 : 0.414, a_3 : 0)$, and $\text{PF}(\mathbf{x}^*, \vec{u}) \leq 1.47$ for all utility profiles $\vec{u} \in \mathcal{U}^{\text{all}}$.⁷

For every utility profile \vec{u} , there exists a distribution \mathbf{x} with $\text{PF}(\mathbf{x}, \vec{u}) = 1$; in fact, the distribution that maximizes Nash welfare with respect to \vec{u} has this property [e.g., Fain et al. 2018, Sec. 2.2]. This is the lowest possible value (no distribution can obtain a value smaller than 1; take $\mathbf{y} = \mathbf{x}$ in the definition of $\text{PF}(\mathbf{x}, \vec{u})$). To illustrate why proportional fairness is a measure of fairness, we can note that if \mathbf{x} is a distribution such that $u_i(\mathbf{x}) = 0$ for some agent $i \in N$, then $\text{PF}(\mathbf{x}, \vec{u}) = \infty$, which we can see by taking any \mathbf{y} for which $u_i(\mathbf{y}) > 0$. Thus, an α -proportionally fair distribution, with α not too high, guarantees to every agent a base level of utility compared to what the agent can receive in any other distribution (in particular, no agent's preferences can be completely ignored).

Like the Nash welfare, proportional fairness is also scale invariant. Hence, we have $\text{PF}_m(f, \mathcal{U}^{\text{all}}) = \text{PF}_m(f, \mathcal{U}^{\text{unit-sum}}) = \text{PF}_m(f, \mathcal{U}^{\text{unit-range}})$ for every voting rule f . Our results for proportional fairness all hold with respect to \mathcal{U}^{all} , so we drop it from the notation and simply use $\text{PF}(\mathbf{x}, \vec{\sigma})$ and $\text{PF}_m(f)$.

⁶This is the maximum possible average multiplicative increase in agent utilities when moving from \mathbf{x} to any other \mathbf{y} . The second transition in Equation (1) holds because $\frac{1}{n} \sum_{i \in N} u_i(\mathbf{y})/u_i(\mathbf{x})$ is linear in \mathbf{y} .

⁷On this small example, one can find this optimum distribution \mathbf{x}^* by hand after noting that a_3 must receive probability 0. One derives $\mathbf{x}^* = \{a_1 : 2 - \sqrt{2}, a_2 : \sqrt{2} - 1, a_3 : 0\}$ with $\text{PF}(\mathbf{x}^*, \vec{\sigma}, \mathcal{U}^{\text{all}}) = 1 + \sqrt{2}/3$.

An appealing strength of proportional fairness is that it is related to other fairness properties of interest. In particular, an α -proportionally fair voting rule is also in the α -core, and has a distortion with respect to Nash welfare of at most α .

Proportional fairness \Rightarrow the core. The following is a well-known relation between proportional fairness and the core.⁸

PROPOSITION 2.5. *For every $\alpha \geq 1$, if f is an α -proportionally fair voting rule, then f is in the α -core.*

PROOF. Suppose for contradiction that $\text{PF}_m(f) \leq \alpha$, but there exists a consistent pair of utility profile \vec{u} and preference profile $\vec{\sigma}$ such that $\mathbf{x} = f(\vec{\sigma})$ is not in the α -core with respect to \vec{u} . Then, by definition, there exists a subset of agents S and a distribution over alternatives $\mathbf{y} \in \Delta(A)$ such that $\frac{|S|}{n} \cdot u_i(\mathbf{y}) \geq \alpha \cdot u_i(\mathbf{x})$ (i.e., $\frac{u_i(\mathbf{y})}{u_i(\mathbf{x})} \geq \alpha \cdot \frac{n}{|S|}$) for every agent $i \in S$ and at least one of these inequalities is strict. Hence,

$$\sum_{i \in S} \frac{u_i(\mathbf{y})}{u_i(\mathbf{x})} > \alpha \cdot n \quad \Rightarrow \quad \frac{1}{n} \sum_{i \in N} \frac{u_i(\mathbf{y})}{u_i(\mathbf{x})} \geq \frac{1}{n} \sum_{i \in S} \frac{u_i(\mathbf{y})}{u_i(\mathbf{x})} > \alpha,$$

contradicting the assumption that $\text{PF}_m(f) \leq \alpha$. □

Proportional fairness \Rightarrow distortion with respect to the Nash welfare. It is also well-known that proportional fairness is an upper bound on the approximation of (i.e., distortion with respect to) the Nash welfare.⁹

PROPOSITION 2.6. *For every voting rule f , we have $D_m^{\text{NW}}(f, \mathcal{U}^{\text{all}}) \leq \text{PF}_m(f, \mathcal{U}^{\text{all}})$.*

PROOF. This holds because for any pair of distributions over alternatives $\mathbf{x}, \mathbf{y} \in \Delta(A)$ and utility profile \vec{u} , we have

$$\frac{\text{NW}(\mathbf{y}, \vec{u})}{\text{NW}(\mathbf{x}, \vec{u})} = \left(\prod_{i \in N} \frac{u_i(\mathbf{y})}{u_i(\mathbf{x})} \right)^{1/n} \leq \frac{1}{n} \sum_{i \in N} \frac{u_i(\mathbf{y})}{u_i(\mathbf{x})},$$

by the inequality of arithmetic and geometric means. □

2.5 The Minimax Theorem

In several places, we will use some basic elements of the theory of zero-sum games and the minimax theorem. Recall that if $X \subseteq \mathbb{R}^n$ is a convex set and $f : X \rightarrow \mathbb{R}$ is a function, then f is *convex* if for all $\mathbf{x}_1, \mathbf{x}_2 \in X$ and all $0 \leq \lambda \leq 1$, we have $f(\lambda \mathbf{x}_1 + (1 - \lambda) \mathbf{x}_2) \leq \lambda f(\mathbf{x}_1) + (1 - \lambda) f(\mathbf{x}_2)$. Further, f is *concave* if $-f$ is convex. For example, $f(x) = x^2$ is convex and $f(x) = \log x$ is concave; linear functions are both convex and concave.

THEOREM 2.7 (MINIMAX THEOREM, VON NEUMANN 1928). *Let $X \subseteq \mathbb{R}^n$ and $Y \subseteq \mathbb{R}^m$ be compact convex sets. Let $f : X \times Y \rightarrow \mathbb{R}$ be a continuous function that is concave in its first argument and convex in its second argument (that is, $f(\cdot, \mathbf{y})$ is concave for each fixed $\mathbf{y} \in Y$ and $f(\mathbf{x}, \cdot)$ is convex for each fixed $\mathbf{x} \in X$). Then*

$$\max_{\mathbf{x} \in X} \min_{\mathbf{y} \in Y} f(\mathbf{x}, \mathbf{y}) = \min_{\mathbf{y} \in Y} \max_{\mathbf{x} \in X} f(\mathbf{x}, \mathbf{y}).$$

⁸This result has not been explicitly stated, but essentially the same proof is frequently used to show that distributions maximizing Nash welfare lie in the core [e.g., Fain et al. 2018, Section 2.2; Aziz et al. 2019, Theorem 3].

⁹Observations to this effect can be found, for example, in Appendix D of Caragiannis et al. [2019] and in the derivation of Equation (2) in Inoue and Kobayashi [2022].

We can interpret this theorem as a statement about a two-player zero-sum game between a *player* and an *adversary*. The player can choose a strategy \mathbf{x} from the set X while aiming to maximize the value $f(\mathbf{x}, \mathbf{y})$, and the adversary can choose $\mathbf{y} \in Y$ aiming to minimize the value. The minimax theorem states that (under certain convexity conditions) it does not matter in which order the players make their moves. In our applications, we have $X = \Delta(S_1)$ and $Y = \Delta(S_2)$ for some finite sets S_1 and S_2 of *pure strategies*, so that X and Y are sets of *mixed strategies*. In this case, the function f typically encodes an expected payoff, $f(\mathbf{x}, \mathbf{y}) = \mathbb{E}_{s_1 \sim \mathbf{x}, s_2 \sim \mathbf{y}}[g(s_1, s_2)]$ for some $g : S_1 \times S_2 \rightarrow \mathbb{R}$. Such an f is linear in both arguments and hence satisfies the conditions of the minimax theorem. In our results about proportional fairness, we will need the full strength of the minimax theorem, allowing for functions f that are not linear in both arguments. The equal value of the max-min and min-max expressions is known as the *value* of the zero-sum game.

3 DISTORTION

We begin by aiming to achieve low distortion with respect to utilitarian social welfare. Boutilier et al. [2015] consider unit-sum utilities and show that any rule must incur distortion at least $\Omega(\sqrt{m})$. They also construct an intricate and artificial voting rule that achieves distortion $O(\sqrt{m} \log^* m)$, thus leaving a tiny gap. They also present a more natural voting rule that achieves distortion $O(\sqrt{m} \log m)$, which we call the *harmonic rule* f_{HR} . It is based on the harmonic scoring rule, according to which each agent i gives $1/k$ points to the alternative she ranks in the k -th position. Given a preference profile $\vec{\sigma}$, the *harmonic score* of an alternative a is $\text{hsc}(a) = \sum_{i \in N} 1/\sigma_i(a)$. Now, with probability $\frac{1}{2}$, the harmonic rule chooses an alternative uniformly at random, and with probability $\frac{1}{2}$, it chooses an alternative a with probability proportional to $\text{hsc}(a)$. Note that the harmonic scores of all alternatives sum to nH_m where $H_m = 1 + \frac{1}{2} + \dots + \frac{1}{m}$. Thus if $\mathbf{x} = f_{\text{HR}}(\vec{\sigma})$ then

$$\mathbf{x}(a) = \frac{1}{2m} + \frac{\text{hsc}(a)}{2nH_m} \text{ for all } a \in A.$$

In this section, we introduce a new rule that achieves distortion $O(\sqrt{m})$, which is optimal up to a constant factor. This rule is based on concepts from cooperative game theory and from the theory of committee selection, and can be computed in polynomial time. Our rule turns out to have robustly good performance, in that its distortion remains low for other utility classes and other welfare functions. We will compare it throughout to the harmonic rule.

3.1 Stable Lotteries

Below the hood, our new voting rule is based on *multi-winner voting*, also known as *committee selection*, which concerns the well-studied problem of selecting a committee $X \subseteq A$ of k alternatives, based on the agents' preferences over the alternatives [Faliszewski et al. 2017]. One aim of the literature on multi-winner voting is to identify *representative* committees, where as many agents as possible are represented in the committee, in the sense that one of their highly-ranked alternatives is included [Chamberlin and Courant 1983]. This is a type of fairness consideration and related to the idea of proportional representation which is particularly well-developed in the context of approval utilities [Lackner and Skowron 2020].

Representative committees are interesting in the distortion context due to the following intuition: if a voting rule places very little weight on alternatives that are highly ranked by many agents, then the rule is in danger of incurring high distortion, because those unrepresented agents may feel strongly about their high-ranked alternatives, while others may be more-or-less indifferent.

For ranked preferences, a recently studied representation axiom is (*local*) *stability* [Aziz et al. 2017b; Cheng et al. 2020]. This axiom is based on the idea that a group of $\frac{n}{k}$ agents should be able to decide over one of the k slots in the committee. Formally, for a committee X with $|X| = k$ and

an alternative a^* , write $V(a^*, X) = |\{i \in N : a^* \succ_i X\}|$ for the number of agents who prefer a^* to all alternatives in the committee. We say that X is *stable* if for all alternatives $a^* \notin X$, we have $V(a^*, X) < \frac{n}{k}$. Such a committee is stable in a sense familiar from cooperative game theory.

There are examples of preference profiles and sizes k where no stable committee exists [Jiang et al. 2020, Thm. 4]. However, Cheng et al. [2020] proved that there always exists a probability distribution over committees which satisfies a probabilistic generalization of the stability property.

Definition 3.1. A distribution $\mathbf{X} \in \Delta(\mathcal{P}_k(A))$ over committees X of size k is a *stable lottery* if for all alternatives $a^* \in A$, we have

$$\mathbb{E}_{X \sim \mathbf{X}} [V(a^*, X)] < \frac{n}{k}.$$

To be self-contained, we include a short proof of existence, following the simplified treatment due to Jiang et al. [2020, Lem. 4].

THEOREM 3.2 (CHENG ET AL. 2020). *For every preference profile $\vec{\sigma}$ and for every k , there exists a stable lottery.*

PROOF. Let $\vec{\sigma}$ be a preference profile. We view our task as proving the following bound:

$$\min_{X \in \Delta(\mathcal{P}_k(A))} \max_{a^* \in A} \mathbb{E}_{X \sim X} [V(a^*, X)] < \frac{n}{k}.$$

If the bound holds, then taking an X that solves the minimization problem is a stable lottery. We can view the expression on the left-hand side as a zero-sum game, where one player chooses a distribution and the adversary responds with an alternative. Applying the minimax theorem, it suffices to show that

$$\max_{\mathbf{y} \in \Delta(A)} \min_{X \in \Delta(\mathcal{P}_k(A))} \mathbb{E}_{X \sim X, a^* \sim \mathbf{y}} [V(a^*, X)] < \frac{n}{k}.$$

Let $\mathbf{y} \in \Delta(A)$. Define a distribution \mathbf{X} over committees by the following process. Draw k alternatives a_1, \dots, a_k from the distribution \mathbf{y} independently and with replacement. Let X be the random set of alternatives thus selected, if necessary filled up with arbitrary additional alternatives until $|X| = k$. Now note that for every agent $i \in N$, the probability $\Pr_{a^* \sim \mathbf{y}, X \sim X} [a^* \succ_i X]$ is at most the probability that a^* is the strictly most-preferred among the at most $k+1$ alternatives a^*, a_1, \dots, a_k which are drawn i.i.d. Hence by symmetry $\Pr_{a^* \sim \mathbf{y}, X \sim X} [a^* \succ_i X] \leq 1/(k+1) < 1/k$.

Summing up over all $i \in N$, it follows that $\mathbb{E}_{X \sim X, a^* \sim \mathbf{y}} [V(a^*, X)] < \frac{n}{k}$, as desired. \square

Cheng et al. [2020] prove that a stable lottery can be found in (expected) polynomial time using the multiplicative weights update algorithm for zero-sum games. That algorithm finds a solution whose value is ε -close to the optimum value. But the existence proof above in fact established that the value of the game is at most $n/(k+1)$, when all we need is a solution with value less than n/k . Thus, we can run the algorithm with $\varepsilon = \frac{1}{2} \cdot (\frac{n}{k} - \frac{n}{k+1})$ and obtain an exactly stable lottery in expected polynomial time.

3.2 The Stable Lottery Rule

We propose a voting rule based on stable lotteries for committees of size $k = \sqrt{m}$. Like the previously proposed harmonic rule, our rule spreads half of the probability mass uniformly over all alternatives.¹⁰ It then assigns the remaining probability mass to alternatives in proportion to the probability that they are included in the committee selected by the stable lottery.

¹⁰Instead of $1/2$, one can use any other constant fraction (such as 0.0001) without changing the conclusion of Theorem 3.4. One can also shift probability from a Pareto-dominated alternatives to a dominating alternative without worsening distortion.

Definition 3.3 (Stable Lottery Rule, f_{SLR}). Let \mathbf{X} be a stable lottery over committees X of size $k = \sqrt{m}$. The Stable Lottery Rule f_{SLR} works as follows: With probability $1/2$, sample a committee $X \sim \mathbf{X}$ and choose an alternative uniformly at random from X , and with probability $1/2$, choose an alternative uniformly at random from A . Therefore, each alternative $a \in A$ will be selected with probability $\mathbf{x}(a) = \frac{1}{2\sqrt{m}} \cdot \Pr_{X \sim \mathbf{X}}[a \in X] + \frac{1}{2m}$.

Our first main result states that f_{SLR} achieves distortion $\Theta(\sqrt{m})$ on the class of balanced utility functions, and hence also for unit-sum, unit-range, and approval utilities.

THEOREM 3.4. *On the utility class $\mathcal{U}^{\text{balanced}}$, the Stable Lottery Rule achieves $O(\sqrt{m})$ distortion:*

$$D_m(f_{\text{SLR}}, \mathcal{U}^{\text{balanced}}) = O(\sqrt{m}).$$

PROOF. Let \vec{u} be a utility profile consistent with a profile $\vec{\sigma}$, with $u_i \in \mathcal{U}^{\text{balanced}}$ for all $i \in N$. We begin the proof by making the following observation. Let X be a committee, and let $a^* \in A$ be a distinguished alternative. Write $u_i(X) = \sum_{a \in X} u_i(a)$ and $\text{UW}(X, \vec{u}) = \sum_{i \in N} u_i(X)$. Then,

$$\text{UW}(X, \vec{u}) \geq \text{UW}(a^*, \vec{u}) - V(a^*, X). \quad (2)$$

Indeed, for every agent i such that $a^* \succ_i X$, we have $u_i(X) \geq 0 \geq u_i(a^*) - \max_{a \in A} u_i(a) \geq u_i(a^*) - 1$ because $u_i \in \mathcal{U}^{\text{balanced}}$, and for every agent i such that $a^* \not\succeq_i X$, there exists some alternative $a \in X$ such that $a \succ_i a^*$, so $u_i(X) \geq u_i(a) \geq u_i(a^*)$. Equation (2) follows by summing these inequalities over all $i \in N$, noting that the number of agents of the first type is $V(a^*, X)$.

Let $\mathbf{x} = f_{\text{SLR}}(\vec{\sigma})$ be the distribution selected the Stable Lottery Rule, and let \mathbf{X} be the underlying stable lottery over committees of size \sqrt{m} . Let us write $\mathbf{x} = \frac{1}{2}\mathbf{x}_1 + \frac{1}{2}\mathbf{x}_2$, where \mathbf{x}_1 is the part of \mathbf{x} based on the stable lottery and \mathbf{x}_2 is the uniform distribution over A . Thus, $\mathbf{x}_1(a) = \frac{1}{\sqrt{m}} \cdot \Pr_{X \sim \mathbf{X}}[a \in X]$ and $\mathbf{x}_2(a) = 1/m$ for all $a \in A$.

Note that for all $i \in N$, we have $u_i(\mathbf{x}_2) \geq \frac{1}{m} \sum_{a \in A} u_i(a) \geq \frac{1}{m}$ because $u_i \in \mathcal{U}^{\text{balanced}}$. Hence $\text{UW}(\mathbf{x}_2, \vec{u}) \geq \frac{n}{m}$ and so $\frac{m}{n} \cdot \text{UW}(\mathbf{x}_2, \vec{u}) \geq 1$. Now fix an arbitrary $a^* \in A$. Then,

$$\begin{aligned} \sqrt{m} \cdot \text{UW}(\mathbf{x}_1, \vec{u}) &= \sqrt{m} \cdot \sum_{a \in A} \frac{1}{\sqrt{m}} \Pr_{X \sim \mathbf{X}}[a \in X] \cdot \text{UW}(a, \vec{u}) \\ &= \mathbb{E}_{X \sim \mathbf{X}}[\sum_{a \in X} \text{UW}(a, \vec{u})] \\ &= \mathbb{E}_{X \sim \mathbf{X}}[\text{UW}(X, \vec{u})] \\ &\geq \mathbb{E}_{X \sim \mathbf{X}}[\text{UW}(a^*, \vec{u}) - V(a^*, X)] && \text{(by equation (2))} \\ &= \text{UW}(a^*, \vec{u}) - \mathbb{E}_{X \sim \mathbf{X}}[V(a^*, X)] && \text{(linearity of expectation)} \\ &\geq \text{UW}(a^*, \vec{u}) - \frac{n}{\sqrt{m}} && \text{(stability of } \mathbf{X} \text{)} \\ &\geq \text{UW}(a^*, \vec{u}) - \frac{n}{\sqrt{m}} \cdot \frac{m}{n} \cdot \text{UW}(\mathbf{x}_2, \vec{u}) && \text{(since } \frac{m}{n} \cdot \text{UW}(\mathbf{x}_2, \vec{u}) \geq 1 \text{)} \\ &= \text{UW}(a^*, \vec{u}) - \sqrt{m} \cdot \text{UW}(\mathbf{x}_2, \vec{u}). \end{aligned}$$

Hence,

$$\text{UW}(\mathbf{x}, \vec{u}) = \frac{1}{2} \text{UW}(\mathbf{x}_1, \vec{u}) + \frac{1}{2} \text{UW}(\mathbf{x}_2, \vec{u}) \geq \frac{\text{UW}(a^*, \vec{u})}{2\sqrt{m}} \quad \text{for all } a^* \in A.$$

Therefore, we have that

$$D_m(f_{\text{SLR}}, \mathcal{U}^{\text{balanced}}) \leq D(\mathbf{x}, \vec{\sigma}, \mathcal{U}^{\text{balanced}}) \leq \max_{a^* \in A} \frac{\text{UW}(a^*, \vec{u})}{\frac{\text{UW}(a^*, \vec{u})}{2\sqrt{m}}} = 2\sqrt{m} = O(\sqrt{m}). \quad \square$$

This proof generalizes to utility functions that are imbalanced to some degree.

COROLLARY 3.5. For $0 < \alpha < 1$, write $\mathcal{U}^{\alpha\text{-balanced}}$ for the class of α -balanced utility functions with $\alpha \cdot \max_{a \in A} u(a) \leq 1 \leq \sum_{a \in A} u(a)$. On this class, the Stable Lottery Rule achieves $O(\sqrt{m/\alpha})$ distortion:

$$D_m(f_{\text{SLR}}, \mathcal{U}^{\alpha\text{-balanced}}) = O(\sqrt{m/\alpha}).$$

In contrast to the $O(\sqrt{m})$ distortion of the Stable Lottery Rule, the Harmonic Rule f_{HR} achieves worse distortion for both unit-sum and, especially, unit-range utilities.

THEOREM 3.6. The distortion of the Harmonic Rule is $D_m(f_{\text{HR}}, \mathcal{U}^{\text{unit-sum}}) = \Theta(\sqrt{m \log m})$ and $D_m(f_{\text{HR}}, \mathcal{U}^{\text{unit-range}}) = \Theta(m^{2/3} \log^{1/3} m)$.

For unit-sum utilities, the upper bound is due to Boutilier et al. [2015] and the lower bound follows from the work of Bhaskar et al. [2018], though we include an explicit lower bound example in Appendix C.1. The analysis of the distortion of f_{HR} for unit-range utilities is new. The polynomial increase in the distortion of f_{HR} compared to that of f_{SLR} can be explained by noting that f_{HR} is truthful, and for unit-range utilities, Filos-Ratsikas and Miltersen [2014] prove that any truthful rule has distortion $\Omega(m^{2/3})$, meaning that f_{HR} still has close to the best distortion achievable via truthful rules. We give proofs of these results in Appendix C.1.

3.3 Lower Bounds

Boutilier et al. [2015] prove that the distortion of every voting rule for the class $\mathcal{U}^{\text{unit-sum}}$ of unit-sum utilities is $\Omega(\sqrt{m})$, showing that f_{SLR} is asymptotically optimal on this class. In Appendix A, we generalize their bound slightly to a broader parametric class of utility functions, which in particular implies that the distortion of f_{SLR} on the class of α -balanced utility functions discussed in Corollary 3.5 is optimal in both m and α .

Here, we present a lower bound for the class of approval utility functions, which also applies to the larger class of unit-range utilities. This bound implies that f_{SLR} achieves asymptotically optimal distortion on both of these utility classes.

THEOREM 3.7. For any voting rule f , we have $D_m(f, \mathcal{U}^{\text{approval}}) = \Omega(\sqrt{m})$ and $D_m(f, \mathcal{U}^{\text{unit-range}}) = \Omega(\sqrt{m})$.

PROOF. Assume \sqrt{n} is a positive integer, and let $m = n + \sqrt{n}$. Each agent i ranks alternative a_i first, alternative $a_{n+\lceil i/\sqrt{n} \rceil}$ second, and the remaining alternatives in an arbitrary order. Note that this naturally divides the agents into \sqrt{n} groups, $N_1, \dots, N_{\sqrt{n}}$, where, for $r \in [\sqrt{n}]$, N_r denotes the group of agents who rank alternative a_{n+r} second.

Let f be a voting rule and let \mathbf{x} be the distribution selected by f on this profile. By the pigeonhole principle, there must exist one index $r \in [\sqrt{n}]$ such that $\mathbf{x}(a_{n+r}) \leq 1/\sqrt{n}$. Without loss of generality, assume that $\mathbf{x}(a_{n+1}) \leq 1/\sqrt{n}$. Consider the approval utility profile \vec{u} under which all agents in N_1 approve their top two alternatives (i.e., their top choice and a_{n+1}), and all other agents approve only their top alternative.

Then, $\text{UW}(a_{n+1}, \vec{u}) = \sqrt{n}$ whereas $\text{UW}(a', \vec{u}) = 1$ for every alternative $a' \in A \setminus \{a_{n+1}\}$. Therefore, we have

$$D_m(f, \mathcal{U}^{\text{approval}}) \geq \frac{\text{UW}(a_{n+1}, \vec{u})}{\text{UW}(\mathbf{x}, \vec{u})} \geq \frac{\sqrt{n}}{(1 - \frac{1}{\sqrt{n}}) \cdot 1 + \frac{1}{\sqrt{n}} \cdot \sqrt{n}} \geq \frac{\sqrt{n}}{2} = \Omega(\sqrt{m}),$$

where the final transition holds due to $m = n + \sqrt{n}$. Because $\mathcal{U}^{\text{approval}} \subseteq \mathcal{U}^{\text{unit-range}}$, we also have $D_m(f, \mathcal{U}^{\text{unit-range}}) = \Omega(\sqrt{m})$. \square

4 PROPORTIONAL FAIRNESS

In this section, we turn our attention to proportional fairness (see Definition 2.3). As we noted in Section 2.4, the proportional fairness objective is scale invariant, and thus $\text{PF}_m(f, \mathcal{U}^{\text{all}}) = \text{PF}_m(f, \mathcal{U}^{\text{unit-sum}}) = \text{PF}_m(f, \mathcal{U}^{\text{unit-range}})$ for all voting rules f . Thus, we will just consider \mathcal{U}^{all} throughout this section, and thus suppress the utility class \mathcal{U} from our notation.

4.1 Upper Bounds

A natural question at this point is whether the stable-lottery-based approach from the previous section, which provides optimal distortion, also works for proportional fairness. In Appendix E, we present a close cousin of our stable lottery rule, namely the *stable committee rule* (f_{SCR}), which uses an approximately stable *deterministic* committee in place of an exactly stable *lottery* over committees; such committees with constant approximations are guaranteed to exist due to the recent work of Jiang et al. [2020]. In the appendix, we show that this rule is $O(\sqrt{m})$ -proportionally fair. This raises the obvious question of whether it is possible to do better. Surprisingly, we show that it is! Using the minimax theorem, we are able to show that there exists an $O(\log m)$ -proportionally fair voting rule. We later show this upper bound to be tight. In Section 4.3, we use the projected subgradient descent algorithm to turn this non-constructive argument into an efficient algorithm.

We begin with a useful lemma that simplifies the analysis of the proportional fairness of a given distribution \mathbf{x} . Let us write $h_i(a) = \{a' \in A : a' \succ_i a\}$ for the set of alternatives that agent i ranks weakly above a , and for a distribution \mathbf{x} , let $\mathbf{x}(h_i(a)) = \sum_{a' \in h_i(a)} \mathbf{x}(a')$ be the total weight that \mathbf{x} places on them.

LEMMA 4.1. *Given a preference profile $\vec{\sigma}$ and a distribution \mathbf{x} , we have*

$$\text{PF}(\mathbf{x}, \vec{\sigma}) = \max_{a \in A} \frac{1}{n} \sum_{i \in N} \frac{1}{\mathbf{x}(h_i(a))}, \quad (3)$$

and this is convex in \mathbf{x} .

PROOF. Recall from Section 2.4 that

$$\text{PF}(\mathbf{x}, \vec{\sigma}) = \sup_{\vec{u} \in (\mathcal{U}^{\text{all}})^n : \vec{u} \succ \vec{\sigma}} \max_{a \in A} \frac{1}{n} \cdot \sum_{i \in N} \frac{u_i(a)}{u_i(\mathbf{x})}.$$

Fix any $a \in A$. Note that we can take the worst case over the utility function u_i of each agent i separately as its contribution to the above expression, for any fixed \mathbf{x} and a , is independent of that of the other utility functions.

Thus, it is sufficient to show that $\sup_{u_i \in \mathcal{U}^{\text{all}} : u_i \succ \sigma_i} u_i(a)/u_i(\mathbf{x}) = 1/\mathbf{x}(h_i(a))$. This follows from the simple observation that $u_i(a') \geq u_i(a)$ for all $a' \in h_i(a)$, which implies $u_i(\mathbf{x}) \geq \mathbf{x}(h_i(a)) \cdot u_i(a)$, i.e., $u_i(a)/u_i(\mathbf{x}) \leq 1/\mathbf{x}(h_i(a))$, and noting that this upper bound is achieved by setting, for example, $u_i(a') = 1$ for all $a' \in h_i(a)$ and $u_i(a') = 0$ for all $a' \in A \setminus h_i(a)$.

Convexity in \mathbf{x} follows because the function $g(z) = 1/z$ is a convex function, and taking the sum and maximum of a collection of convex functions yields a convex function. \square

Note that the last line of the proof shows that the worst case for proportional fairness is achieved at an approval utility profile. Hence, $\text{PF}_m(f, \mathcal{U}^{\text{all}}) = \text{PF}_m(f, \mathcal{U}^{\text{approval}})$ for all voting rules f .

With this simplified formulation in hand, we can now derive an upper bound on the optimal proportional fairness.

THEOREM 4.2. *There exists a voting rule which is $2(1 + \ln(2m))$ -proportionally fair.*

PROOF. We consider the instance-optimal voting rule which, given a preference profile, selects a distribution \mathbf{x} that is α -proportionally fair for the smallest α . We interpret this distribution as the

outcome of a (two-player) zero-sum game and α as the value of that game. We then bound this value in a dual game obtained by applying the minimax theorem.

Formulation as a zero-sum game. Let $\vec{\sigma}$ be any preference profile. Write $\text{PF}(\vec{\sigma}) = \min_{\mathbf{x} \in \Delta(A)} \text{PF}(\mathbf{x}, \vec{\sigma})$. Lemma 4.1 implies that

$$\text{PF}(\vec{\sigma}) = \min_{\mathbf{x} \in \Delta(A)} \max_{a \in A} \frac{1}{n} \sum_{i \in N} \frac{1}{\mathbf{x}(h_i(a))}.$$

Hence, $\text{PF}(\vec{\sigma})$ can be viewed as the outcome of a zero-sum game. The set of *pure* strategies for the (first) player is $\Delta(A)$, i.e., the player may choose a distribution over alternatives. In response, the adversary (the second player) can choose a single alternative $a \in A$. Then, for a pair of strategies $(\mathbf{x}, a) \in \Delta(A) \times A$, the payoff to the adversary, which is equal to the negative payoff of the player, is defined as

$$R(\mathbf{x}, a) = \frac{1}{n} \sum_{i \in N} \frac{1}{\mathbf{x}(h_i(a))}.$$

With this notation, we have

$$\text{PF}(\vec{\sigma}) = \min_{\mathbf{x} \in \Delta(A)} \max_{a \in A} R(\mathbf{x}, a).$$

Suppose we allow the adversary to choose a *mixed* strategy, i.e., a distribution over alternatives $S_A \in \Delta(A)$. Define the expected payoff of the pair (\mathbf{x}, S_A) of strategies to be $\mathbb{E}_{a \sim S_A} [R(\mathbf{x}, a)]$. Because this objective is linear in S_A , there is always a pure best response for the adversary (selecting a single alternative $a \in A$). Thus, allowing the adversary to choose a mixed strategy does not change the value of the game. Hence

$$\text{PF}(\vec{\sigma}) = \min_{\mathbf{x} \in \Delta(A)} \max_{S_A \in \Delta(A)} \mathbb{E}_{a \sim S_A} [R(\mathbf{x}, a)].$$

Now note that $\mathbb{E}_{a \sim S_A} [R(\mathbf{x}, a)]$ is convex in \mathbf{x} (Lemma 4.1) and linear (and hence concave) in S_A . Therefore, by the minimax theorem (Theorem 2.7), we have

$$\text{PF}(\vec{\sigma}) = \max_{S_A \in \Delta(A)} \min_{\mathbf{x} \in \Delta(A)} \mathbb{E}_{a \sim S_A} [R(\mathbf{x}, a)].$$

We call this game the *dual game*.

Bounding the value of the dual game. In the dual game, for a given strategy S_A of the adversary, suppose the player responds with the strategy $\bar{\mathbf{x}}$ with $\bar{\mathbf{x}}(a) = \frac{1}{2}S_A(a) + \frac{1}{2m}$ for all $a \in A$ (which is not necessarily a best response). Thus, with probability $\frac{1}{2}$ the player selects according to S_A , and with probability $\frac{1}{2}$, the player selects an alternative uniformly at random. Note that the value of the dual game when the player plays $\bar{\mathbf{x}}$ is an upper bound on the true value of the dual game. Now, we have

$$\begin{aligned} \text{PF}(\vec{\sigma}) &= \max_{S_A \in \Delta(A)} \min_{\mathbf{x} \in \Delta(A)} \mathbb{E}_{a \sim S_A} [R(\mathbf{x}, a)] \\ &\leq \max_{S_A \in \Delta(A)} \mathbb{E}_{a \sim S_A} [R(\bar{\mathbf{x}}, a)] && \text{(first player responds with } \bar{\mathbf{x}} \text{)} \\ &= \max_{S_A \in \Delta(A)} \frac{1}{n} \sum_{i \in N} \mathbb{E}_{a \sim S_A} \left[\frac{1}{\bar{\mathbf{x}}(h_i(a))} \right] && \text{(linearity of expectation)} \\ &\leq \max_{S_A \in \Delta(A)} \max_{i \in N} \mathbb{E}_{a \sim S_A} \left[\frac{1}{\bar{\mathbf{x}}(h_i(a))} \right]. \end{aligned}$$

The last term is maximized at some distribution S_A and some agent i with preference ranking σ_i . Without loss of generality, suppose $\sigma_i = a_1 \succ_i a_2 \succ_i \dots \succ_i a_m$. Write $T_j = \sum_{\ell=1}^j \bar{\mathbf{x}}(a_\ell)$ and $T_0 = 0$.

Furthermore, note that, in this case, $T_j = \bar{x}(h_i(a_j))$. Then the above is equal to

$$\sum_{j \in [m]} S_\Lambda(a_j) \cdot \frac{1}{\bar{x}(h_i(a_j))} < \sum_{j \in [m]} \frac{2 \cdot \bar{x}(a_j)}{\bar{x}(h_i(a_j))} = 2 \sum_{j \in [m]} \frac{T_j - T_{j-1}}{T_j} = 2 \sum_{j \in [m]} \left(1 - \frac{T_{j-1}}{T_j}\right).$$

Using the fact that $1 - x \leq -\ln(x)$,

$$\sum_{j \in [m]} \left(1 - \frac{T_{j-1}}{T_j}\right) \leq 1 + \sum_{j=2}^m (\ln(T_j) - \ln(T_{j-1})) = 1 + \ln(T_m) - \ln(T_1) \leq 1 + \ln(2m),$$

where the last inequality holds due to $T_m = 1$ and $T_1 = \mathbf{x}(a_1) \geq \frac{1}{2m}$. It follows that $\text{PF}(\vec{\sigma}) \leq 2(1 + \ln(2m))$, as desired. \square

This upper bound on proportional fairness immediately implies upper bounds on the α -core and distortion with respect to Nash welfare, using Propositions 2.5 and 2.6.

COROLLARY 4.3. *Let $\alpha = 2(1 + \ln(2m))$. There exists a voting rule f which is in the α -core and whose distortion with respect to Nash welfare is $D_m^{\text{NW}}(f) \leq \alpha$.*

We can also show that the harmonic rule f_{HR} is $\Theta(\sqrt{m \log m})$ -proportionally fair. While this is worse by a polynomial factor compared to our $O(\log m)$ -proportionally fair rule, it is close to optimal for truthful rules, which we show cannot be better than $\Omega(\sqrt{m})$ -proportionally fair.

THEOREM 4.4. *The harmonic rule f_{HR} is $\Theta(\sqrt{m \log m})$ -proportionally fair.*

THEOREM 4.5. *If f is an α -proportionally fair voting rule that is also truthful, then $\alpha = \Omega(\sqrt{m})$.*

We provide proofs of these results in Appendix C.2 and Appendix D, respectively.

4.2 Lower Bounds

Next, we give a lower bound that essentially matches our upper bound, showing it is optimal.

THEOREM 4.6. *If f is an α -proportionally fair voting rule, then $\alpha \geq H_m \geq \ln(m+1)$.*

PROOF. We show that there exists a preference profile $\vec{\sigma}$ for which $\text{PF}(\mathbf{x}, \vec{\sigma}) \geq \ln(m+1)$ for every distribution $\mathbf{x} \in \Delta(A)$; this implies the desired claim.

For simplicity, we assume $n = m$, but the proof works when n is a multiple of m . Take the cyclic preference profile $\vec{\sigma}$, in which each agent $i \in [m]$ has the preference ranking $\sigma_i = a_i > a_{i+1} > \dots > a_m > a_1 > \dots > a_{i-1}$. Let \mathbf{x} be any distribution. For $i \in [m]$ and $\ell \in [m]$, define $T_i(\ell) = \sum_{k=0}^{\ell-1} \mathbf{x}(a_{i+k})$ with the subscripts cycling as $a_{m+k} \triangleq a_k$ for all $k \in [m]$. Using this notation, we have $\mathbf{x}(h_i(a)) = T_i(\sigma_i(a))$ for any $i \in [m]$ and $a \in A$. Then,

$$\begin{aligned} \text{PF}(\mathbf{x}, \vec{\sigma}) &= \max_{a \in A} \frac{1}{m} \sum_{i \in [m]} \frac{1}{T_i(\sigma_i(a))} \geq \frac{1}{m^2} \sum_{a \in A} \sum_{i \in [m]} \frac{1}{T_i(\sigma_i(a))} \\ &= \frac{1}{m^2} \sum_{\ell \in [m]} \sum_{i \in [m]} \frac{1}{T_i(\ell)} \geq \sum_{\ell \in [m]} \frac{1}{\sum_{i \in [m]} T_i(\ell)}, \end{aligned}$$

where the final transition is the standard inequality between arithmetic and harmonic means.

For each fixed ℓ , note that $\sum_{i \in [m]} T_i(\ell) = \ell$ as each alternative appears in the top ℓ positions in exactly ℓ of the rankings. Using this above, we get $\text{PF}(\mathbf{x}, \vec{\sigma}) \geq \sum_{\ell \in [m]} 1/\ell = H_m \geq \ln(m+1)$. \square

The profile used in the proof of Theorem 4.6 can also be used to derive a lower bound on distortion with respect to Nash welfare using a similar analysis. We provide the proof in Appendix F.

PROPOSITION 4.7. *Every voting rule has distortion at least $(\frac{m^m}{m!})^{\frac{1}{m}}$ with respect to the Nash welfare.*

It can be checked that the lower bound from Proposition 4.7 converges to e from below as $m \rightarrow \infty$. For $m = 2$, the lower bound is $\sqrt{2}$. In Appendix G, we show that this is tight by proposing a voting rule that achieves the $\sqrt{2}$ bound.

4.3 Computation

Lemma 4.1 gives a simple formula for calculating the value $\text{PF}(\mathbf{x}, \vec{\sigma})$ of a given distribution \mathbf{x} . Now, we turn to the computational problem of finding an α -proportionally fair distribution \mathbf{x} with the lowest possible α , for a given preference profile $\vec{\sigma}$. We show that this can be (approximately) achieved in polynomial time. Our argument depends on the convexity of $\text{PF}(\mathbf{x}, \vec{\sigma})$ in \mathbf{x} (Lemma 4.1), which allows us to use convex optimization methods (in particular, the projected subgradient descent algorithm). For definitions of a subgradient and the subdifferential ∂f of a convex function f , we refer the reader to the books of Nesterov [2003]; Vishnoi [2021].

THEOREM 4.8 (NESTEROV [2003], CHAPTER 3.2, VISHNOI [2021], THEOREM 7.1). *Let f be a convex function over a bounded closed convex set Q . There is an algorithm (based on subgradient descent) that, given (a) an oracle that, given $\mathbf{x} \in Q$, can return $f(\mathbf{x})$ and a subgradient $\mathbf{g} \in \partial f(\mathbf{x})$; (b) a number G such that for all $\mathbf{x} \in Q$ and subgradients $\mathbf{g} \in \partial f(\mathbf{x})$, we have $\|\mathbf{g}\|_2 \leq G$; (c) an initial point $\mathbf{x}^0 \in Q$; (d) a number D such that $\|\mathbf{x}^0 - \mathbf{x}^*\|_2 \leq D$ where $\mathbf{x}^* = \arg \min_{\mathbf{x} \in Q} f(\mathbf{x})$; and (e) some $\varepsilon > 0$, outputs a sequence $\mathbf{x}^0, \mathbf{x}^1, \dots, \mathbf{x}^{T-1}$ such that $\frac{1}{T} \sum_{i=0}^{T-1} f(\mathbf{x}^i) - f(\mathbf{x}^*) \leq \varepsilon$, where $T = \left(\frac{DG}{\varepsilon}\right)^2$.*

As in the proof of Theorem 4.2, let $\text{PF}(\vec{\sigma}) = \min_{\mathbf{x} \in \Delta(A)} \text{PF}(\mathbf{x}, \vec{\sigma})$ denote the best possible approximation of proportional fairness on a given preference profile $\vec{\sigma}$.

THEOREM 4.9. *Given a preference profile $\vec{\sigma}$ and $\varepsilon > 0$, an α -proportionally fair distribution over alternatives with $\alpha \leq \text{PF}(\vec{\sigma}) + \varepsilon$ can be computed in $\text{poly}(n, m, 1/\varepsilon)$ time.*

PROOF. We apply the projected subgradient method described in Theorem 4.8 to $f(\mathbf{x}) = \text{PF}(\mathbf{x}, \vec{\sigma}) = \max_{a \in A} R(\mathbf{x}, a)$, where $R(\mathbf{x}, a) = \frac{1}{n} \sum_{i \in N} 1/x(h_i(a))$. We have shown f to be convex in \mathbf{x} (Lemma 4.1). Note that for a given $\mathbf{x} \in \Delta(A)$, $f(\mathbf{x})$ can be computed in $\text{poly}(n, m)$ time. We want to optimize f over $\mathbf{x} \in \Delta(A)$. However, the subgradients at points close to the boundary of the probability simplex can be unbounded. We will avoid this by carefully restricting the domain of f .

For each $a \in A$, let p_a be the fraction of agents who rank a as their top choice. Let $\beta = 2(1 + \ln(2m))$ be the upper bound proven in Theorem 4.2 on $\text{PF}(\vec{\sigma})$. We set $Q = \{\mathbf{x} \in \Delta(A) : \mathbf{x}(a) \geq p_a/\beta, \forall a \in A\}$.

First, we will show that an optimal distribution $\mathbf{x}^* \in \arg \max_{\mathbf{x} \in \Delta(A)} f(\mathbf{x})$ lies in Q , ensuring that it is sufficient to optimize f over Q . Subsequently, we will show that the norm of any subgradient of f at any point in Q is bounded by $\text{poly}(n, m)$ and that such a subgradient can be computed in polynomial time, giving us conditions (a) and (b) of Theorem 4.8. Then, the only missing piece left to be able to apply Theorem 4.8 is a starting point: we can choose any $\mathbf{x}^0 \in Q$ (e.g., $\mathbf{x}^0(a) = p_a$ for each $a \in A$) and use $D = \sqrt{2}$, which is an upper bound on the Euclidean distance between any two probability distributions.

Optimality. Let $\mathbf{x}^* \in \arg \min_{\mathbf{x} \in \Delta(A)} f(\mathbf{x})$. Assume for a contradiction that $\mathbf{x}^* \notin Q$. Thus, there exists an alternative $a \in A$ with $p_a > 0$ such that $\mathbf{x}^*(a) < p_a/\beta$. Then,

$$\text{PF}(\vec{\sigma}) = \text{PF}(\mathbf{x}^*, \vec{\sigma}) \geq \frac{1}{n} \sum_{i \in N} \frac{1}{\mathbf{x}^*(h_i(a))} \geq p_a \cdot \frac{1}{\mathbf{x}^*(a)} > \beta.$$

This contradicts Theorem 4.2, where we proved that $\text{PF}(\vec{\sigma}) \leq \beta$ for all $\vec{\sigma}$. Therefore, $\mathbf{x}^* \in Q$.

Bounding and computing subgradients. Take any $\mathbf{x} \in Q$. For each fixed $a \in A$, the function $R(\mathbf{x}, a)$ is differentiable and convex in \mathbf{x} . More specifically, for all $a, a' \in A$, we have

$$\frac{\partial}{\partial x_{a'}} R(\mathbf{x}, a) = \frac{1}{n} \sum_{i \in N: a' \succ_i a} \frac{-1}{\mathbf{x}(h_i(a))^2}.$$

For each $i \in N$, let a_i^* be the top alternative of agent i . Because $p_{a_i^*} \geq 1/n$ and $\mathbf{x} \in Q$, we have

$$\mathbf{x}(h_i(a)) \geq \mathbf{x}(a_i^*) \geq \frac{p_{a_i^*}}{\beta} \geq \frac{1}{2n(1 + \ln(2m))}.$$

Therefore, $\|\nabla R(\mathbf{x}, a)\|_\infty = O((n \ln m)^2)$ for all a . Furthermore, it is known that any gradient $\nabla R(\mathbf{x}, a^*)$, where $a^* \in \arg \max_{a \in A} R(\mathbf{x}, a)$ is a subgradient of $\text{PF}(\mathbf{x}, \vec{\sigma}) = \max_{a \in A} R(\mathbf{x}, a)$. Hence, it follows that the norm of such a subgradient is bounded by $\text{poly}(n, m)$ and we can compute such a subgradient in $\text{poly}(n, m)$ time. \square

5 DISCUSSION

We have proved that the best distortion (with respect to the utilitarian welfare) that probabilistic voting rules can achieve with ranked preferences is $\Theta(\sqrt{m})$, resolving an open question by Boutilier et al. [2015]. We have also initiated the study of proportional fairness as the counterpart of distortion which focuses on fairness rather than efficiency, proving that the value of this objective with ranked preferences is $\Theta(\log m)$.

The $O(\log m)$ upper bound on the approximation of proportional fairness carries over to distortion with respect to the Nash welfare. However, our lower bound for this objective is at most e , leaving open the question of whether constant distortion with respect to the Nash welfare is achievable. Similarly, one can also focus on distortion with respect to other welfare functions, such as the egalitarian welfare or, more generally, the p -mean welfare [Barman et al. 2020; Chaudhury et al. 2021]. For the egalitarian welfare, it is easy to see that the best distortion for approval utilities is $\Theta(m)$,¹¹ but it is unclear what the answer is for unit-sum utilities.

Beyond our setting, there is significant literature on studying distortion with respect to the utilitarian welfare for ballot formats other than ranked preferences [Amanatidis et al. 2021; Benade et al. 2021; Borodin et al. 2022; Mandal et al. 2019, 2020]. A natural direction for future work is to study proportional fairness and distortion with respect to other welfare functions for such ballot formats. One can also extend these ideas from single-winner selection to committee selection, where the output of a voting rule is a (randomized) subset of alternatives of a given size, and participatory budgeting, where each alternative has a cost and the output is a (randomized) subset of alternatives with total cost at most a given budget.

Finally, centuries of research on voting theory has focused on *simple* voting rules (such as plurality or Borda count) that are easy for voters to understand and satisfy appealing axiomatic properties. A significant barrier to the modern optimization-based approaches, which focus on quantitative objectives such as distortion or proportional fairness, is that they often yield rules that are difficult to understand (and sometimes difficult to compute). Significant challenges lie ahead in paving the path for increased practicability of such approaches: Can we design simple rules that perform well on these quantitative metrics? Alternatively, can we convey the intricate rules emerging from such approaches to the end users by providing simple-to-digest *explanations* of either their end goal or their properties [Peters et al. 2021b]? Can we reconcile these quantitative approaches with the classical axiomatic approach to find rules that achieve the best of both worlds?

¹¹The upper bound can be achieved by assigning a probability of $1/m$ to each alternative. For the lower bound, consider a cyclic profile over m alternatives. Any probabilistic voting rule assigns probability at most $1/m$ to some alternative; one can set approval utilities such that the egalitarian welfare is 1 for this alternative and 0 for every other alternative.

REFERENCES

- [1] Stéphane Airiau, Haris Aziz, Ioannis Caragiannis, Justin Kruger, Jérôme Lang, and Dominik Peters. 2019. Portioning using Ordinal Preferences: Fairness and Efficiency. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*. 11–17.
- [2] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, and Alexandros A. Voudouris. 2021. Peeking behind the ordinal curtain: Improving distortion via cardinal queries. *Artificial Intelligence* 296 (2021), 103488.
- [3] Elliot Anshelevich, Onkar Bhardwaj, Edith Elkind, John Postl, and Piotr Skowron. 2018. Approximating optimal social choice under metric preferences. *Artificial Intelligence* 264 (2018), 27–51.
- [4] Elliot Anshelevich and John Postl. 2017. Randomized social choice functions under metric preferences. *Journal of Artificial Intelligence Research* 58 (2017), 797–827.
- [5] Elliot Anshelevich and Shreyas Sekar. 2016. Blind, greedy, and random: Algorithms for matching and clustering using only ordinal information. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*. 383–389.
- [6] Haris Aziz, Anna Bogomolnaia, and Hervé Moulin. 2019. Fair mixing: the case of dichotomous preferences. In *Proceedings of the 20th ACM Conference on Economics and Computation (EC)*. 753–781.
- [7] Haris Aziz, Florian Brandl, Felix Brandt, and Markus Brill. 2018. On the Tradeoff between Efficiency and Strategyproofness. *Games and Economic Behavior* 110 (2018), 1–18.
- [8] Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh. 2017. Justified representation in approval-based committee voting. *Social Choice and Welfare* 48, 2 (2017), 461–485.
- [9] Haris Aziz, Edith Elkind, Piotr Faliszewski, Martin Lackner, and Piotr Skowron. 2017. The Condorcet principle for multiwinner elections: from shortlisting to proportionality. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*. 84–90.
- [10] Salvador Barberà. 1978. Nice decision schemes. In *Decision Theory and Social Ethics*. 101–117.
- [11] Siddharth Barman, Umang Bhaskar, Anand Krishna, and Ranjani G. Sundaram. 2020. Tight Approximation Algorithms for p -Mean Welfare Under Subadditive Valuations. In *Proceedings of the 28th Annual European Symposium on Algorithms (ESA)*. 11:1–11:17.
- [12] Gerdus Benade, Swaprava Nath, Ariel D. Procaccia, and Nisarg Shah. 2021. Preference elicitation for participatory budgeting. *Management Science* 67, 5 (2021), 2813–2827.
- [13] Jeremy Bentham. 1789. *An Introduction to the Principles of Morals and Legislation*. T. Payne, London.
- [14] Umang Bhaskar, Varsha Dani, and Abheek Ghosh. 2018. Truthful and near-optimal mechanisms for welfare maximization in multi-winner elections. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*. 925–932.
- [15] Anna Bogomolnaia, Hervé Moulin, and Richard Stong. 2005. Collective choice under dichotomous preferences. *Journal of Economic Theory* 122, 2 (2005), 165–184.
- [16] Allan Borodin, Daniel Halpern, Mohamad Latifian, and Nisarg Shah. 2022. Distortion in Voting with Top- t Preferences. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*. forthcoming.
- [17] Craig Boutilier, Ioannis Caragiannis, Simi Haber, Tyler Lu, Ariel D. Procaccia, and Or Sheffet. 2015. Optimal Social Choice Functions: A Utilitarian View. *Artificial Intelligence* 227 (2015), 190–213.
- [18] Florian Brandl, Felix Brandt, Dominik Peters, and Christian Stricker. 2021. Distribution Rules Under Dichotomous Preferences: Two Out of Three Ain’t Bad. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*. 158–179.
- [19] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. 2019. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Transactions on Economics and Computation* 7, 3 (2019).
- [20] Ioannis Caragiannis, Swaprava Nath, Ariel D. Procaccia, and Nisarg Shah. 2017. Subset Selection Via Implicit Utilitarian Voting. *Journal of Artificial Intelligence Research* 58 (2017), 123–152.
- [21] Ioannis Caragiannis and Ariel D Procaccia. 2011. Voting almost maximizes social welfare despite limited communication. *Artificial Intelligence* 175, 9-10 (2011), 1655–1671.
- [22] Ioannis Caragiannis, Nisarg Shah, and Alexandros A. Voudouris. 2022. The Metric Distortion of Multiwinner Voting. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*. Forthcoming.
- [23] John R. Chamberlin and Paul N. Courant. 1983. Representative Deliberations and Representative Decisions: Proportional Representation and the Borda Rule. *The American Political Science Review* 77, 3 (1983), 718–733.
- [24] Moses Charikar and Prasanna Ramakrishnan. 2022. Metric Distortion Bounds for Randomized Social Choice. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2986–3004.
- [25] Bhaskar Ray Chaudhury, Jugal Garg, and Ruta Mehta. 2021. Fair and Efficient Allocations under Subadditive Valuations. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*. 5269–5276.
- [26] Yu Cheng, Zhihao Jiang, Kamesh Munagala, and Kangning Wang. 2020. Group fairness in committee selection. *ACM Transactions on Economics and Computation (TEAC)* 8, 4 (2020), 1–18.
- [27] Conal Duddy. 2015. Fair sharing under dichotomous preferences. *Mathematical Social Sciences* 73 (2015), 1–5.

- [28] Brandon Fain, Ashish Goel, and Kamesh Munagala. 2016. The core of the participatory budgeting problem. In *Proceedings of the 12th Conference on Web and Internet Economics (WINE)*. 384–399.
- [29] Brandon Fain, Kamesh Munagala, and Nisarg Shah. 2018. Fair Allocation of Indivisible Public Goods. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*. 575–592. Full version arXiv:1805.03164.
- [30] Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. 2017. Multiwinner Voting: A New Challenge for Social Choice Theory. In *Trends in Computational Social Choice*, Ulle Endriss (Ed.), Chapter 2.
- [31] Aris Filos-Ratsikas. 2015. *Social Welfare in Algorithmic Mechanism Design Without Money*. Ph.D. Dissertation. Aarhus University.
- [32] Aris Filos-Ratsikas, Evi Micha, and Alexandros A. Voudouris. 2020. The distortion of distributed voting. *Artificial Intelligence* 286 (2020), 103343.
- [33] Aris Filos-Ratsikas and Peter Bro Miltersen. 2014. Truthful approximations to range voting. In *Proceedings of the 10th Conference on Web and Internet Economics (WINE)*. 175–188.
- [34] Vasilis Gkatzelis, Daniel Halpern, and Nisarg Shah. 2020. Resolving the optimal metric distortion conjecture. In *Proceedings of the 61st Annual Symposium on Foundations of Computer Science (FOCS)*. 1427–1438.
- [35] Asei Inoue and Yusuke Kobayashi. 2022. An Additive Approximation Scheme for the Nash Social Welfare Maximization with Identical Additive Valuations. *arXiv:2201.01419* (2022).
- [36] Zhihao Jiang, Kamesh Munagala, and Kangning Wang. 2020. Approximately stable committee selection. In *Proceedings of the 52nd Annual ACM Symposium on Theory of Computing (STOC)*. 463–472.
- [37] Frank P. Kelly, Aman K. Maulloo, and David K. H. Tan. 1998. Rate control for communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society* 49, 3 (1998), 237–252.
- [38] David Kempe. 2020. Communication, distortion, and randomness in metric voting. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*. 2087–2094.
- [39] Martin Lackner and Piotr Skowron. 2020. Approval-Based Committee Voting. (2020). arXiv:2007.01795.
- [40] Sinya Lee. 2013. Maximization of Relative Social Welfare on Truthful Cardinal Voting Schemes. *arXiv:1904.00538* (2013).
- [41] Debmalya Mandal, Ariel D. Procaccia, Nisarg Shah, and David P. Woodruff. 2019. Efficient and Thrifty Voting by Any Means Necessary. In *Proceedings of the 33rd Annual Conference on Neural Information Processing Systems (NeurIPS)*. 7180–7191.
- [42] Debmalya Mandal, Nisarg Shah, and David P. Woodruff. 2020. Optimal Communication-Distortion Tradeoff in Voting. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*. 795–813.
- [43] Kamesh Munagala, Yiheng Shen, Kangning Wang, and Zhiyi Wang. 2022. Approximate Core for Committee Selection via Multilinear Extension and Market Clearing. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2229–2252.
- [44] Kamesh Munagala and Kangning Wang. 2019. Improved metric distortion for deterministic social choice rules. In *Proceedings of the 20th ACM Conference on Economics and Computation (EC)*. 245–262.
- [45] Yurii Nesterov. 2003. *Introductory Lectures on Convex Optimization*. Springer.
- [46] Dominik Peters, Grzegorz Pierczyński, Nisarg Shah, and Piotr Skowron. 2021. Market-based explanations of collective decisions. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*. 5656–5663.
- [47] Dominik Peters, Grzegorz Pierczyński, and Piotr Skowron. 2021. Proportional Participatory Budgeting with Additive Utilities. In *Proceedings of the 34th Annual Conference on Neural Information Processing Systems (NeurIPS)*. 12726–12737.
- [48] Dominik Peters and Piotr Skowron. 2020. Proportionality and the limits of welfarism. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*. 793–794.
- [49] Ariel D. Procaccia and Jeffrey S. Rosenschein. 2006. The distortion of cardinal preferences in voting. In *Proceedings of the 10th International Workshop on Cooperative Information Agents (CIA)*. Springer, 317–331.
- [50] Luis Sánchez-Fernández, Edith Elkind, Martin Lackner, Norberto Fernández, Jesús Fisteus, Pablo Basanta Val, and Piotr Skowron. 2017. Proportional justified representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*. 670–676.
- [51] Piotr Skowron. 2021. Proportionality degree of multiwinner rules. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*. 820–840.
- [52] Nisheeth K. Vishnoi. 2021. *Algorithms for Convex Optimization*. Cambridge University Press.
- [53] John von Neumann. 1928. Zur Theorie der Gesellschaftspiele. *Math. Ann.* 100, 1 (1928), 295–320.

APPENDIX

A LOWER BOUND FOR ALMOST BALANCED UTILITIES

Here, we generalize the lower bound of Boutilier et al. [17] to give a lower bound on the distortion of the utility class of utilities that are at most slightly imbalanced.

For $\varepsilon \in [1/m, 1]$, let \mathcal{U}^ε be the class containing the following utility functions:

- For each $a^* \in A$, it contains a utility function u under which $u(a^*) = 1$ and $u(a) = 0$ for all $a \in A \setminus \{a^*\}$. Note that this corresponds to “single-minded” agents.
- It contains a utility function u under which $u(a) = \varepsilon/m$ for all $a \in A$. Note that this corresponds to agents who are indifferent among all alternatives.

We prove that every voting rule has distortion $\Omega(\sqrt{m/\varepsilon})$ on \mathcal{U}^ε . Because $\mathcal{U}^{\text{unit-sum}} \supset \mathcal{U}^{\varepsilon=1}$ and because the distortion (weakly) increases when the utility class grows, this implies $\Omega(\sqrt{m})$ bound for $\mathcal{U}^{\text{unit-sum}}$, generalizing the result of Boutilier et al. [17]. We also note that $\mathcal{U}^{\text{approval}} \supset \mathcal{U}^{\varepsilon=m}$; however, the construction below does not immediately work for $\varepsilon > 1$. We leave it open to check if it can be generalized to larger ε , perhaps in a way that subsumes Theorem 3.7 (our lower bound on distortion for approval utilities) as a special case. Finally, we note that $\mathcal{U}^\varepsilon \subseteq \mathcal{U}^{\varepsilon\text{-balanced}}$, implying that the distortion of f_{SLR} on the class of α -balanced utilities discussed in Corollary 3.5 is optimal.

THEOREM A.1. *For any $\varepsilon \in [1/m, 1]$, the distortion of every voting rule f on the utility class \mathcal{U}^ε satisfies $D_m(f, \mathcal{U}) = \Omega(\sqrt{m/\varepsilon})$.*

PROOF. Assume $\sqrt{m/\varepsilon}$ is a positive integer. Fix a set of alternatives $T = \{a_1, \dots, a_{\sqrt{m/\varepsilon}}\}$, and partition the set of agents into $\sqrt{m/\varepsilon}$ buckets $B_1, \dots, B_{\sqrt{m/\varepsilon}}$, where each bucket consists of $n/\sqrt{m/\varepsilon}$ agents. Now, construct a profile $\vec{\sigma}$ in which for all $j \in [\sqrt{m/\varepsilon}]$, all agents in B_j rank alternative a_j first, and the remaining alternatives are ranked arbitrarily.

Let f be a voting rule, and let \mathbf{x} be the distribution that f selects on this profile. By the pigeonhole principle, there must exist an index $i \in [\sqrt{m/\varepsilon}]$ such that $\mathbf{x}(a_i) \leq \frac{1}{\sqrt{m/\varepsilon}}$. Now, construct a utility profile \vec{u} where all agents in B_i assign utility 1 to a_i and utility 0 to all other alternatives, and where all agents in other blocks assign utility ε/m to all alternatives. We can see that $\text{UW}(a_i, \vec{u}) \geq \frac{n}{\sqrt{m/\varepsilon}}$ from the agents in B_i , whereas $\text{UW}(a_j, \vec{u}) \leq n \cdot \frac{\varepsilon}{m}$ for all $j \neq i$. This also means that a_i is the alternative with maximum utility in this profile.

Now, we have

$$\text{UW}(\mathbf{x}, \vec{u}) = \mathbf{x}(a_i) \cdot \text{UW}(a_i, \vec{u}) + \sum_{j \neq i} \mathbf{x}(a_j) \cdot \text{UW}(a_j, \vec{u}) \leq \frac{1}{\sqrt{m/\varepsilon}} \cdot \frac{n}{\sqrt{m/\varepsilon}} + 1 \cdot \frac{n\varepsilon}{m} = \frac{2n\varepsilon}{m}.$$

This implies the distortion of f on \mathcal{U}^ε satisfies

$$D_m(f, \mathcal{U}^\varepsilon) \geq \frac{\text{UW}(a_i, \vec{u})}{\text{UW}(\mathbf{x}, \vec{u})} \geq \frac{\frac{n}{\sqrt{m/\varepsilon}}}{\frac{2n\varepsilon}{m}} = \frac{1}{2} \cdot \sqrt{\frac{m}{\varepsilon}} = \Omega\left(\sqrt{\frac{m}{\varepsilon}}\right),$$

as needed. □

B APPROVAL VS. UNIT-RANGE UTILITIES

In this section, we show that approval utilities are the worst case for distortion among the broader class of unit-range utilities. Hence, any upper bounds derived for approval utilities (like in Theorem C.2 in the next section) apply to the broader class of unit-range utilities as well. The proof also implies that a distribution \mathbf{x} that minimizes $D(\mathbf{x}, \vec{\sigma}, \mathcal{U}^{\text{unit-range}})$ thereby also minimizes

$D(\mathbf{x}, \vec{\sigma}, \mathcal{U}^{\text{approval}})$, which simplifies the computational problem of optimizing the latter quantity (see the discussion after Example 2.2).

LEMMA B.1. *For every voting rule f , we have $D_m(f, \mathcal{U}^{\text{unit-range}}) = D_m(f, \mathcal{U}^{\text{approval}})$.*

PROOF. As $\mathcal{U}^{\text{approval}} \subseteq \mathcal{U}^{\text{unit-range}}$, we trivially have $D_m(f, \mathcal{U}^{\text{approval}}) \leq D_m(f, \mathcal{U}^{\text{unit-range}})$. We show that the inequality also holds in the opposite direction. We will prove a stronger argument: for every distribution \mathbf{x} and preference profile $\vec{\sigma}$, we have $D(\mathbf{x}, \vec{\sigma}, \mathcal{U}^{\text{unit-range}}) \leq D(\mathbf{x}, \vec{\sigma}, \mathcal{U}^{\text{approval}})$. Fix any distribution \mathbf{x} and preference profile $\vec{\sigma}$.

Let $\vec{u} \in (\mathcal{U}^{\text{unit-range}})^n$ be a utility profile consistent with $\vec{\sigma}$ that maximizes $D(\mathbf{x}, \vec{u})$, and among all such utility profiles, let it be the one that minimizes the number of agents who do not have approval utilities. If $\vec{u} \in (\mathcal{U}^{\text{approval}})^n$, then we are done. Suppose this is not the case. Fix any agent i such that $u_i \notin \mathcal{U}^{\text{approval}}$.

Let $a^* \in \arg \max_{a \in A} \text{UW}(a, \vec{u})$ be an alternative maximizing utilitarian welfare under \vec{u} . Then,

$$D(\mathbf{x}, \vec{u}) = \frac{\text{UW}(a^*, \vec{u})}{\text{UW}(\mathbf{x}, \vec{u})} = \frac{\text{UW}(a^*, \vec{u}_{-i}) + u_i(a^*)}{\text{UW}(\mathbf{x}, \vec{u}_{-i}) + \sum_{a \in A} \mathbf{x}(a) \cdot u_i(a)}, \quad (4)$$

where \vec{u}_{-i} denotes the utility profile containing the utility functions of all agents except agent i .

If a^* is the top alternative of agent i (i.e., $\sigma_i(a^*) = 1$), then by the definition of unit-range utilities, we must have $u_i(a^*) = 1$. In that case, it is easy to see that the expression in Equation (4) is maximized when $u_i(a) = 0$ for all $a \in A \setminus \{a^*\}$. That is, define \vec{u}^* such that $\vec{u}_{-i}^* = \vec{u}_{-i}$, $u_i^*(a^*) = 1$, and $u_i^*(a) = 0$ for all $a \in A \setminus \{a^*\}$. Then, $D(\mathbf{x}, \vec{u}^*) \geq D(\mathbf{x}, \vec{u})$, which is a contradiction because \vec{u}^* has at least one more agent with an approval utility function compared to \vec{u} .

Now, suppose $\sigma_i(a^*) \geq 2$. Denote by a^+ the top alternative of agent i satisfying $\sigma_i(a^+) = 1$. Write $\tilde{h}_i(a^*) = \{a \in A \setminus \{a^+\} : a \succsim_i a^*\}$. Consider a different utility profile \vec{u}' , where $\vec{u}'_{-i} = \vec{u}_{-i}$, $u'_i(a^+) = u_i(a^+) = 1$, $u'_i(a) = u_i(a^*)$ for all $a \in \tilde{h}_i(a^*)$, and $u_i(a) = 0$ for all $a \in A$ with $a^* \succ_i a$. Note that we are reducing the utility of agent i for any alternative she ranks higher than a^* to $u_i(a^*)$, and reducing her utility for any alternative she ranks lower than a^* to 0, without changing her utility for a^* . This can only (weakly) reduce the denominator in Equation (4) without changing the numerator, implying that $D(\mathbf{x}, \vec{u}') \geq D(\mathbf{x}, \vec{u})$.

Next, notice that

$$\begin{aligned} D(\mathbf{x}, \vec{u}') &= \frac{\text{UW}(a^*, \vec{u}_{-i}) + u_i(a^*)}{\text{UW}(\mathbf{x}, \vec{u}_{-i}) + \mathbf{x}(a^+) \cdot 1 + \mathbf{x}(\tilde{h}_i(a^*)) \cdot u_i(a^*)} \\ &\leq \max \left(\frac{\text{UW}(a^*, \vec{u}_{-i}) + 1}{\text{UW}(\mathbf{x}, \vec{u}_{-i}) + \mathbf{x}(a^+) \cdot 1 + \mathbf{x}(\tilde{h}_i(a^*)) \cdot 1}, \frac{\text{UW}(a^*, \vec{u}_{-i})}{\text{UW}(\mathbf{x}, \vec{u}_{-i}) + \mathbf{x}(a^+) \cdot 1} \right), \end{aligned}$$

where the final transition holds due to the (weighted) mediant inequality.¹²

Hence, we can see that one of two choices – either increasing the utility of agent i for all the alternatives in $\tilde{h}_i(a^*)$ to 1 or decreasing them all to 0 – does not reduce the distortion. This yields another utility profile \vec{u}^* consistent with $\vec{\sigma}$ such that $D(\mathbf{x}, \vec{u}^*) \geq D(\mathbf{x}, \vec{u}') \geq D(\mathbf{x}, \vec{u})$, but \vec{u}^* has at least one more agent having an approval utility function, which is the desired contradiction. \square

In the proof of Lemma 4.1, we proved that the same conclusion holds for proportional fairness. Because proportional fairness is scale-invariant, this actually holds with respect to the class \mathcal{U}^{all} of all utility functions.

LEMMA B.2. *For every voting rule f , we have $\text{PF}_m(f, \mathcal{U}^{\text{all}}) = \text{PF}_m(f, \mathcal{U}^{\text{approval}})$.*

¹²The general form is given as $\min_i \frac{a_i}{b_i} \leq \frac{\sum_i w_i a_i}{\sum_i w_i b_i} \leq \max_i \frac{a_i}{b_i}$. Here, we take $(a_1, b_1) = (\text{UW}(a^*, \vec{u}_{-i}), \text{UW}(\mathbf{x}, \vec{u}_{-i}) + \mathbf{x}(a^+))$, $(a_2, b_2) = (\text{UW}(a^*, \vec{u}_{-i}) + 1, \text{UW}(\mathbf{x}, \vec{u}_{-i}) + \mathbf{x}(a^+) + \mathbf{x}(\tilde{h}_i(a^*)))$, and $(w_1, w_2) = (1 - u_i(a^*), u_i(a^*))$.

Using a slight generalization of that argument, one can prove this for distortion with respect to Nash welfare too.

LEMMA B.3. *For every voting rule f , we have $D^{\text{NW}}(f, \mathcal{U}^{\text{all}}) = D^{\text{NW}}(f, \mathcal{U}^{\text{approval}})$.*

PROOF. We prove something stronger: for every distribution \mathbf{x} and preference profile $\vec{\sigma}$, we have $D^{\text{NW}}(\mathbf{x}, \vec{\sigma}, \mathcal{U}^{\text{all}}) = D^{\text{NW}}(\mathbf{x}, \vec{\sigma}, \mathcal{U}^{\text{approval}})$. Recall that

$$D^{\text{NW}}(\mathbf{x}, \vec{\sigma}, \mathcal{U}^{\text{all}}) = \sup_{\vec{u} \in (\mathcal{U}^{\text{all}})^n: \vec{u} \succ \vec{\sigma}} \sup_y \left(\prod_{i \in N} \frac{u_i(\mathbf{y})}{u_i(\mathbf{x})} \right)^{1/n}.$$

First, as in the proof of Lemma 4.1, we see that we can take the worst case over the utility function u_i of each agent i separately as its contribution to the distortion expression is independent of that of the other utility functions. Thus, it is sufficient to prove that for fixed distributions \mathbf{x} and \mathbf{y} and agent i , there is an approval utility function u_i that maximizes $\frac{u_i(\mathbf{y})}{u_i(\mathbf{x})}$ across all utility functions consistent with $\vec{\sigma}$. Fix any distributions \mathbf{x} and \mathbf{y} , agent i , and utility function $u_i : A \rightarrow \mathbb{R}_{\geq 0}$.

For simplicity, label alternatives so that σ_i is $a_1 \succ_i a_2 \succ_i \dots \succ_i a_m$, and hence $u_i(a_1) \geq \dots \geq u_i(a_m)$. Take the m different approval utility functions consistent with σ_i : for all $j \in [m]$, let v_j be the utility function that approves alternatives a_1 to a_j . Note that u_i can be written as a non-negative linear combination of the approval utilities, that is, $u_i = \sum_{j \in [m]} \alpha_j v_j$ for some $\alpha_1, \dots, \alpha_m \geq 0$. (Explicitly, we can take $\alpha_m = u_i(a_m)$ and $\alpha_j = u_i(a_j) - u_i(a_{j+1}) \geq 0$ for each $j < m$.) Then,

$$\frac{u_i(\mathbf{y})}{u_i(\mathbf{x})} = \frac{\sum_{j \in [m]} \alpha_j v_j(\mathbf{y})}{\sum_{j \in [m]} \alpha_j v_j(\mathbf{x})} \leq \max_{j \in [m]} \frac{v_j(\mathbf{y})}{v_j(\mathbf{x})},$$

where the final transition is due to the weighted median inequality given as $\min_k \frac{a_k}{b_k} \leq \frac{\sum_k \alpha_k a_k}{\sum_k \alpha_k b_k} \leq \max_k \frac{a_k}{b_k}$. This proves that $u_i(\mathbf{y})/u_i(\mathbf{x}) \leq v_j(\mathbf{y})/v_j(\mathbf{x})$ for some approval utility function v_j , as desired. \square

C THE HARMONIC RULE

In this section, we provide a detailed analysis of the harmonic rule f_{HR} proposed by Boutilier et al. [17]. Recall that for each $a \in A$, we write $\text{hsc}(a) := \sum_{i \in N} 1/\sigma_i(a)$ for its harmonic score. With probability $1/2$, f_{HR} chooses an alternative uniformly at random, and with probability $1/2$, f_{HR} chooses an alternative proportionally to its harmonic score. In other words, we choose each $a \in A$ with probability $\mathbf{x}(a) := \frac{1}{2m} + \frac{\text{hsc}(a)}{2 \sum_{a' \in A} \text{hsc}(a')}$. Note that $\sum_{a' \in A} \text{hsc}(a') = nH_m$, where $H_m := \sum_{i \in [m]} 1/m$ is the m^{th} harmonic number, so we may rewrite $\mathbf{x}(a) = \frac{1}{2m} + \frac{\text{hsc}(a)}{2nH_m}$.

C.1 Distortion

Boutilier et al. [17] show that the distortion of the harmonic rule for unit-sum utilities is $D_m(f_{\text{HR}}, \mathcal{U}^{\text{unit-sum}}) = O(\sqrt{m \log m})$. Bhaskar et al. [14] show that every truthful rule f incurs distortion $D_m(f, \mathcal{U}^{\text{unit-sum}}) = \Omega(\sqrt{m \log m})$. Since f_{HR} is a truthful rule, this implies that the analysis of Boutilier et al. [17] is tight and the harmonic rule has distortion exactly $\Theta(\sqrt{m \log m})$.

For convenience, we include an explicit proof of the lower bound that does not use truthfulness.

THEOREM C.1. $D_m(f_{\text{HR}}, \mathcal{U}^{\text{unit-sum}}) = \Omega(\sqrt{m \log m})$.

PROOF. Consider the preference profile $\vec{\sigma}$ with $n = m - 1$ agents (the construction also works when n is a multiple of $m - 1$), in which each agent places a distinguished alternative a^* at position $k = \sqrt{\frac{m}{2H_m}}$ (for simplicity, assume this is an integer) and the remaining alternatives are arranged

cyclically in the remaining positions so that every remaining alternative appears in every remaining position once. Consider a consistent utility profile \vec{u} in which each agent has utility $1/k$ for her k most preferred alternatives and utility 0 for all other alternatives.

First, note that the optimal social welfare is $\text{UW}(a^*, \vec{u}) = (m-1) \cdot \frac{1}{k} \geq \frac{m}{2k}$, where the last transition holds for $m \geq 2$. In contrast, for any $a \in A \setminus \{a^*\}$, we have $\text{UW}(a, \vec{u}) = (k-1) \cdot \frac{1}{k} \leq 1$ because alternative a is among the top k alternatives of precisely $k-1$ agents.

Finally, the harmonic score of a^* is $\text{hsc}(a^*, \vec{\sigma}) = (m-1) \cdot 1/k$, meaning that the harmonic rule f_{HR} selects a^* with probability $\mathbf{x}(a^*) = \frac{1}{2m} + \frac{1/k}{2H_m} \leq \max(\frac{1}{m}, \frac{1}{2kH_m})$. Hence, the distortion of f_{HR} satisfies

$$\begin{aligned} D_m(f_{\text{HR}}, \mathcal{U}^{\text{unit-sum}}) &\geq \frac{\text{UW}(a^*, \vec{u})}{\mathbf{x}(a^*) \cdot \text{UW}(a^*, \vec{u}) + (1 - \mathbf{x}(a^*)) \cdot 1} \\ &\geq \frac{\text{UW}(a^*, \vec{u})}{\mathbf{x}(a^*) \cdot \text{UW}(a^*, \vec{u}) + 1} \\ &= \frac{1}{\mathbf{x}(a^*) + \frac{1}{\text{UW}(a^*, \vec{u})}} \\ &\geq \frac{1}{\max\left(\frac{1}{m}, \frac{1}{2kH_m}\right) + \frac{2k}{m}} \\ &\geq \frac{1}{2 \max\left(\frac{1}{m}, \frac{1}{2kH_m}, \frac{2k}{m}\right)}. \end{aligned}$$

Setting $k = \sqrt{\frac{m}{2H_m}}$, we get that the distortion is at least $\min(m/2, \sqrt{mH_m/8}) = \Omega(m \log m)$. \square

Next, we analyze the distortion of f_{HR} for unit-range and approval utilities. Strikingly, while the distortion of f_{HR} for unit-sum utilities is only a sublogarithmic factor worse than the best possible distortion, we find that its distortion for approval and unit-range utilities is $\Theta(m^{2/3} \log^{1/3} m)$, which is worse than the best possible distortion of $\Theta(\sqrt{m})$ for these utility classes by a polynomial factor. This contrast can be explained due to a result of Filos-Ratsikas and Miltersen [33] and Lee [40]: for unit-range utilities, the best distortion achieved by any *strategyproof* voting rule is $\Theta(m^{2/3})$.¹³ Because f_{HR} is known to be strategyproof [14], it must have distortion $\Omega(m^{2/3})$; further, its distortion for unit-range utilities is still only a sublogarithmic factor worse than that of the best *strategyproof* voting rule. Note that our stable lottery rule f_{SLR} achieves $\Theta(\sqrt{m})$ distortion for both unit-sum and approval utilities.

THEOREM C.2. *The distortion of the harmonic rule with respect to the class of unit-range and of approval utilities is $D_m(f_{\text{HR}}, \mathcal{U}^{\text{unit-range}}) = D_m(f_{\text{HR}}, \mathcal{U}^{\text{approval}}) = \Theta(m^{2/3} \log^{1/3} m)$.*

PROOF. We begin by proving the upper bound.

Upper bound. By Lemma B.1, for proving the upper bound for unit-range utilities, it suffices to consider approval utilities. Fix an arbitrary profile $\vec{\sigma}$, and let \vec{u} be some consistent utility profile of approval utilities. For each agent $i \in N$, let $r_i = \sum_{a \in A} u_i(a) \geq 1$ be the number of alternatives approved by i . Let $a^* \in \arg \max_{a \in A} \text{UW}(a, \vec{u})$ be an optimal alternative, and let $\mathbf{x} = f_{\text{HR}}(\vec{\sigma})$ be the distribution selected by the harmonic rule.

Let τ be a threshold value to be set later. Consider two cases.

¹³Strategyproofness is defined in Appendix C.2.

Case 1: Suppose $\text{hsc}(a^*) \geq \tau$. Then $\mathbf{x}(a^*) \geq \frac{1}{2} \cdot \frac{\tau}{n \cdot H_m}$ and so $\text{UW}(\mathbf{x}, \vec{u}) \geq \frac{1}{2} \cdot \frac{\tau}{n \cdot H_m} \text{UW}(a^*, \vec{u})$. Thus $D(\mathbf{x}, \vec{u}) = \text{UW}(a^*, \vec{u}) / \text{UW}(\mathbf{x}, \vec{u}) \leq 2nH_m / \tau$.

Case 2: Suppose $\text{hsc}(a^*) \leq \tau$. Let $Y = \{i \in N : u_i(a^*) = 1\}$ be the set of agents approving a^* . Note that $\text{UW}(a^*, \vec{u}) = |Y|$. Because $\mathbf{x}(a) \geq 1/(2m)$ for each $a \in A$, we have $\text{UW}(\mathbf{x}, \vec{u}) \geq \sum_{i \in N} \frac{r_i}{2m}$. Thus

$$D(\mathbf{x}, \vec{u}) = \text{UW}(a^*, \vec{u}) / \text{UW}(\mathbf{x}, \vec{u}) \leq 2m|Y| / (\sum_{i \in N} r_i). \quad (5)$$

We will upper bound this quantity in two different ways.

First, we have $D(\mathbf{x}, \vec{u}) \leq 2m|Y|/n$ because $r_i \geq 1$ for each $i \in N$.

Second, we can observe that

$$\tau \geq \text{hsc}(a^*) \geq \sum_{i \in Y} \frac{1}{r_i} \geq \frac{|Y|^2}{\sum_{i \in Y} r_i} \geq \frac{|Y|^2}{\sum_{i \in N} r_i}, \quad (6)$$

where the first inequality is due to the assumption of Case 2, the second inequality is because every $i \in Y$ ranks a^* among the first r_i positions, and the third inequality is the AM-HM inequality. Rewriting (6), we have $|Y| / (\sum_{i \in N} r_i) \leq \tau / |Y|$. Plugging this into (5), we see that $D(\mathbf{x}, \vec{u}) \leq 2m\tau / |Y|$.

Combining the two bounds, and using $\min(x, y) \leq \sqrt{xy}$, we see that in Case 2 we have

$$D(\mathbf{x}, \vec{u}) \leq \min\left(\frac{2m|Y|}{n}, \frac{2m\tau}{|Y|}\right) \leq 2m \sqrt{\frac{\tau}{n}}.$$

Finally, combining Case 1 and Case 2, we can see that the distortion is at most

$$D(\mathbf{x}, \vec{u}) \leq \max\left(2nH_m/\tau, 2m \sqrt{\frac{\tau}{n}}\right).$$

Setting $\tau = n \cdot (H_m/m)^{2/3}$ yields the optimal upper bound of $2H_m^{1/3} m^{2/3}$.

Lower bound. Assume $m \geq 2$ without loss of generality. Let $t = (H_m/m)^{1/3}$ and $r = 1/t = (m/H_m)^{1/3}$. Choose an arbitrary alternative $a^* \in A$ and construct a preference profile as follows:

- Alternative a^* is ranked r -th by $n \cdot t$ “special” agents and m -th by the remaining $n \cdot (1 - t)$ “ordinary” agents.
- The remaining preferences are filled arbitrarily subject to the condition that each of the remaining $m - 1$ alternatives appear as the top choice of $n \cdot (1 - t) / (m - 1)$ ordinary agents and in the first $r - 1$ positions in the preference rankings of $n \cdot t \cdot (r - 1) / (m - 1)$ special agents.

We set a consistent utility profile as follows:

- Every special agent has utility 1 for her top r alternatives and 0 for the rest.
- Every ordinary agent has utility 1 for her top alternative and 0 for the rest.

Let us analyze the harmonic scores and welfare of various alternatives. For our chosen alternative a^* , we have

$$\text{hsc}(a^*) = n \cdot \left(\frac{t}{r} + \frac{1-t}{m}\right) \leq n \cdot (t^2 + 1/m) \leq 2nt^2,$$

where the final transition uses the fact that $1/m \leq t^2 = (H_m/m)^{2/3}$. Based on this, we get that the probability of a^* being chosen under f_{HR} is

$$\Pr(a^*) \leq \frac{1}{2} \cdot \frac{2t^2}{H_m} + \frac{1}{2m} \leq \frac{3t^2}{2H_m},$$

where the final transition uses the fact that $1/m \leq t^2/H_m$. Next, the social welfare of a^* is

$$\text{UW}(a^*, \vec{u}) = n \cdot t,$$

whereas the social welfare of every other alternative $a \in A \setminus \{a^*\}$ is

$$\text{UW}(a, \vec{u}) = \frac{n \cdot t \cdot (r-1)}{m-1} + \frac{n \cdot (1-t)}{m-1} \leq \frac{n}{m-1} \cdot (tr+1) = \frac{4n}{m},$$

where the final transition uses $rt = 1$ and $m-1 \geq m/2$.

Hence, we get that the distortion of f_{HR} is at least

$$\frac{n \cdot t}{n \cdot t \cdot \frac{3t^2}{2H_m} + \frac{4n}{m} \cdot 1} = \frac{2 \cdot m \cdot t}{11} = \frac{2}{11} \cdot H_m^{1/3} \cdot m^{2/3},$$

as needed. \square

C.2 Proportional Fairness

Finally, we show that the harmonic rule is $\Theta(\sqrt{m \log m})$ -proportionally fair. However, since we have shown that the best voting rule is $\Theta(\log m)$ -proportionally fair, this is also worse by a polynomial factor. Later, we explain this contrast once again via strategyproofness of the harmonic rule, by proving that if an α -proportionally fair voting rule is strategyproof, then $\alpha = \Omega(\sqrt{m})$. Hence, the harmonic rule is only a sublogarithmic factor worse than the best strategyproof voting rule according to the proportional fairness metric.

THEOREM C.3. *The harmonic rule is $\Theta(\sqrt{m \log m})$ -proportionally fair.*

PROOF. Let us begin by proving the upper bound. Fix a consistent pair of utility profile \vec{u} and preference profile $\vec{\sigma}$. Let $\mathbf{x} = f_{\text{HR}}(\vec{\sigma})$ be the probability distribution returned by f_{HR} on $\vec{\sigma}$. Recall that

$$\text{PF}(\mathbf{x}, \vec{u}) = \max_{a \in A} \frac{1}{n} \cdot \sum_{i \in N} \frac{u_i(a)}{u_i(\mathbf{x})} \leq \max_{a \in A} \frac{1}{n} \cdot \sum_{i \in N} \frac{1}{\mathbf{x}(h_i(a))}, \quad (7)$$

where $h_i(a) = \{b : b \succeq_i a\}$ is the set of alternatives that agent i ranks at least as high as a . Let a^* denote arg max of the right-hand side above.

For $r \in [m]$, let α_r denote the fraction of agents who rank a^* in position r . Note that $\sum_{r=1}^m \alpha_r = 1$. Further, the harmonic score of a^* is given by $\text{hsc}(a^*) = n \cdot \sum_{r=1}^m \alpha_r / r$. We consider two cases.

Case 1. Suppose $\sum_{r=1}^m \frac{\alpha_r}{r} \geq \sqrt{H_m/m}$. Then, $\text{hsc}(a^*) \geq n\sqrt{H_m/m}$. Hence,

$$\mathbf{x}(h_i(a^*)) \geq \mathbf{x}(a^*) \geq \frac{1}{2} \cdot \frac{1}{\sqrt{mH_m}}.$$

Plugging this in Equation (7), we get that $\text{PF}(\mathbf{x}, \vec{u}) \leq 2\sqrt{mH_m}$, as desired.

Case 2. Suppose $\sum_{r=1}^m \frac{\alpha_r}{r} \leq \sqrt{H_m/m}$. Note that $\mathbf{x}(a) \geq \frac{1}{2m}$ for every alternative $a \in A$. Hence, if agent i ranks a^* in position r , we have $\mathbf{x}(h_i(a)) \geq \frac{r}{2m}$. Plugging this into Equation (7), we get

$$\text{PF}(\mathbf{x}, \vec{u}) \leq \sum_{r=1}^m \frac{(2m) \cdot \alpha_r}{r} \leq 2\sqrt{mH_m},$$

as desired.

Next, we prove the lower bound. Fix a special alternative a^* . Construct a preference profile $\vec{\sigma}$ in which there are $n = m-1$ agents. Alternative a^* is ranked in position $r = \sqrt{m/H_m}$ by all the agents. The other alternatives are placed in the remaining positions in a cyclic manner, so that every other

alternative appears in every remaining position exactly once. Let $\mathbf{x} = f_{\text{HR}}(\vec{\sigma})$ be the probability distribution returned by the harmonic rule on this profile. Note that

$$\mathbf{x}(a^*) = \frac{1}{2} \frac{\text{hsc}(a^*)}{nH_m} + \frac{1}{2} \frac{1}{m} = \frac{1}{2} \frac{1}{rH_m} + \frac{1}{2} \frac{1}{m} \leq \frac{1}{\sqrt{mH_m}},$$

where the last inequality holds because $m \geq H_m$. By symmetry, the remaining probability is equally distributed among the remaining alternatives. Hence, we have $\mathbf{x}(a) \leq 1/(m-1)$ for all $a \in A \setminus \{a^*\}$.

Next, fix a utility profile \vec{u} in which every agent i has utility 1 for her r most favorite alternatives. Note that for every agent $i \in N$, we have

$$u_i(\mathbf{x}) \leq \frac{1}{\sqrt{mH_m}} + \frac{r-1}{m-1} \leq \frac{1}{\sqrt{mH_m}} + \frac{r}{m} = \frac{2}{\sqrt{mH_m}}.$$

In contrast, $u_i(a^*) = 1$ for all agents $i \in N$. Hence,

$$\text{PF}(\mathbf{x}, \vec{u}) \geq \frac{1}{n} \sum_{i \in N} \frac{u_i(a^*)}{u_i(\mathbf{x})} \geq \frac{\sqrt{mH_m}}{2},$$

as desired. \square

D PROPORTIONAL FAIRNESS OF STRATEGYPROOF VOTING RULES

In this section, we prove that every strategyproof voting rule can only be $\Omega(\sqrt{m})$ -proportionally fair. Since the harmonic rule is $\Theta(\sqrt{m \log m})$ -proportionally fair, it is at most a sublogarithmic factor worse than the best strategyproof voting rule. Let us now define strategyproofness formally.

Definition D.1 (Strategyproofness). A voting rule f is called *strategyproof* (also known as *truthful*) if no agent can increase her utility by misreporting her vote. Formally, for any preference profile $\vec{\sigma}$, any agent $i \in N$, any utility function u_i consistent with σ_i , and any ranking of the alternatives σ'_i , we must have $u_i(f(\vec{\sigma})) \geq u_i(f((\sigma'_i, \vec{\sigma}_{-i})))$, where $(\sigma'_i, \vec{\sigma}_{-i})$ is the preference profile obtained by replacing the vote of agent i in $\vec{\sigma}$ by σ'_i .

Before proving the result, we need to introduce several other definitions. First, we define two well-known and mild properties of (probabilistic) voting rules.

Definition D.2 (Anonymity). A voting rule f is called *anonymous* if its outcome does not depend on the identities of the agents. Formally, for any preference profile $\vec{\sigma}$ and any permutation of the agents $\pi^N : N \rightarrow N$, we must have $f(\pi^N \circ \vec{\sigma}) = f(\vec{\sigma})$, where $\pi^N \circ \vec{\sigma} = (\sigma_{\pi^N(i)})_{i \in N}$ is the profile obtained by permuting the votes in $\vec{\sigma}$ according to π^N .

Definition D.3 (Neutrality). A voting rule f is called *neutral* if its outcome does not depend on the names of the alternatives. Formally, for any preference profile $\vec{\sigma}$ and any permutation of the alternatives $\pi^A : A \rightarrow A$, we must have $f(\pi^A \circ \vec{\sigma}) = \pi^A \circ f(\vec{\sigma})$, where $\pi^A \circ \vec{\sigma}$ is the profile obtained by permuting the alternatives in each vote σ_i according to π^A , and $\pi^A \circ f(\vec{\sigma})$ is the distribution obtained by permuting the names of the alternatives in $f(\vec{\sigma})$ according to π^A .

Next, we introduce two classes of voting rules, following the work of Barberà [10].

Definition D.4 (Point-Voting Schemes). A voting rule f is called a *point-voting scheme* if there exists a vector $\vec{w} = (w_1, \dots, w_m)$ with $w_1 \geq w_2 \geq \dots \geq w_m \geq 0$ and $\sum_{r \in [m]} w_r = 1$, such that for every preference profile $\vec{\sigma}$, writing $\mathbf{x} = f(\vec{\sigma})$, we have

$$\mathbf{x}(a) = \frac{1}{n} \sum_{i \in N} w_{\sigma_i(a)} \quad \text{for all } a \in A.$$

Informally, a point-voting scheme, parametrized by the vector \vec{w} , resembles the positional scoring rule parametrized by the same score vector \vec{w} , except that the positional scoring rule would choose the alternative a with the highest total score $\sum_{i \in N} w_{\sigma_i(a)}$ whereas the point-voting scheme chooses each alternative a with probability proportional to its score. Another way to view a point-voting scheme is that it chooses an agent uniformly at random and then chooses her r -th ranked alternative with probability w_r , for each $r \in [m]$.

Note that the first half of the harmonic rule f_{HR} , which chooses each alternative with probability proportional to its harmonic score, is a point-voting scheme.¹⁴

Definition D.5 (Supporting-Size Schemes). A voting rule f is called a *supporting-size scheme* if there exists a vector $\vec{z} = (z_0, \dots, z_n)$ with

- $z_n \geq z_{n-1} \geq \dots \geq z_0$, and
- $z_k + z_{n-k} = 1$ for each $k \in [n] \cup \{0\}$

such that for every preference profile $\vec{\sigma}$, writing $\mathbf{x} = f(\vec{\sigma})$, we have

$$\mathbf{x}(a) = \frac{1}{\binom{m}{2}} \sum_{b \in A \setminus \{a\}} z_{n_{a>b}} \quad \text{for all } a \in A, \text{ where } n_{a>b} = |\{i \in N : a \succ_i b\}|.$$

In other words, f chooses a pair of alternatives (a, b) uniformly at random and then chooses each of them with probability obtained from the vector \vec{z} at index that is equal to the number of agents preferring that alternative over the other.

Barberà [10] proved the following result.

PROPOSITION D.6. *A voting rule is anonymous, neutral, and strategyproof if and only if it is a probability mixture of a point-voting scheme and a supporting-size scheme.*

The reason this is useful is that by analyzing the best objective value (distortion or proportional fairness) achievable by any mixture of a point-voting scheme and a supporting-size scheme, we also obtain the best objective value achievable by any anonymous, neutral, and strategyproof voting rule. Could a strategyproof voting rule that violates anonymity and/or neutrality achieve a better objective value? For distortion, Filos-Ratsikas and Miltersen [33] prove that this is not the case, and this observation was used by them and by Bhaskar et al. [14] to derive the aforementioned lower bounds on the distortion of any strategyproof voting rule with respect to the unit-range and unit-sum utility classes, respectively. It is easy to see that the same observation holds for proportional fairness as well.

LEMMA D.7. *For every strategyproof voting rule f , there exists an anonymous, neutral, and strategyproof voting rule f' such that $\text{PF}_m(f') \leq \text{PF}_m(f)$.*

PROOF. Like Filos-Ratsikas and Miltersen [33], we consider a strategyproof voting rule f and construct a voting rule f' which works as follows: given an input preference profile $\vec{\sigma}$, it applies a uniformly random permutation of the agents π^N and an independently chosen uniformly random permutation of alternatives π^A to $\vec{\sigma}$, then applies rule f on the resulting profile $\pi^A \circ \pi^N \circ \vec{\sigma}$, and finally applies the inverse of π^A , denoted $(\pi^A)^{-1}$, on the resulting distribution to revert the change of names of alternatives.

Filos-Ratsikas and Miltersen [33] argue that if f is strategyproof, then f' is anonymous, neutral, and strategyproof. Further, $D(f', \mathcal{U}^{\text{unit-range}}) \leq D_m(f, \mathcal{U}^{\text{unit-range}})$ (and $D(f', \mathcal{U}^{\text{unit-sum}}) \leq D_m(f, \mathcal{U}^{\text{unit-sum}})$ using the same argument). We want to show that $\text{PF}_m(f') \leq \text{PF}_m(f)$ as well. The argument for this is slightly more involved because, unlike the social welfare function, proportional

¹⁴In fact, it is easy to observe that the entire harmonic rule is a point-voting scheme with each agent awarding $1/r + H_m/m$ points to the alternative she ranks in position r , for each $r \in [m]$.

fairness is non-linear. Crucially, we use the fact that $\text{PF}(\mathbf{x}, \vec{\sigma})$ is convex in \mathbf{x} as we observed following Lemma 4.1.

Take any preference profile $\vec{\sigma}$ and let $\mathbf{x} = f'(\vec{\sigma})$. Let Π^N and Π^A denote the set of permutations of agents and alternatives, respectively. Note that

$$\mathbf{x} = \frac{1}{|\Pi^N| \cdot |\Pi^A|} \sum_{\pi^N \in \Pi^N, \pi^A \in \Pi^A} (\pi^A)^{-1} \circ f(\pi^A \circ \pi^N \circ \vec{\sigma}).$$

Hence, we have

$$\begin{aligned} \text{PF}(\mathbf{x}, \vec{\sigma}) &\leq \frac{1}{|\Pi^N| \cdot |\Pi^A|} \sum_{\pi^N \in \Pi^N, \pi^A \in \Pi^A} \text{PF}((\pi^A)^{-1} \circ f(\pi^A \circ \pi^N \circ \vec{\sigma}), \vec{\sigma}) \\ &\leq \max_{\pi^N \in \Pi^N, \pi^A \in \Pi^A} \text{PF}((\pi^A)^{-1} \circ f(\pi^A \circ \pi^N \circ \vec{\sigma}), \vec{\sigma}) \\ &= \max_{\pi^N \in \Pi^N, \pi^A \in \Pi^A} \text{PF}((\pi^A)^{-1} \circ f(\pi^A \circ \pi^N \circ \vec{\sigma}), \pi^N \circ \vec{\sigma}) \\ &= \max_{\pi^N \in \Pi^N, \pi^A \in \Pi^A} \text{PF}(f(\pi^A \circ \pi^N \circ \vec{\sigma}), \pi^A \circ \pi^N \circ \vec{\sigma}) \\ &\leq \max_{\vec{\sigma}'} \text{PF}(f(\vec{\sigma}'), \vec{\sigma}') = \text{PF}_m(f), \end{aligned}$$

where the first transition is due to convexity of the proportional fairness objective, the second transition upper bounds an average by the maximum, the third transition uses the fact that proportional fairness is an anonymous objective (i.e., permuting the votes does not change the proportional fairness value of a distribution on the preference profile), the fourth transition uses the fact that proportional fairness is a neutral objective (i.e., permuting the names of alternatives in both the distribution and the preference profile keeps the proportional fairness value unchanged), and the final transition upper bounds the maximum over a subset of preference profiles by a maximum over all preference profiles. Since this holds for each preference profile $\vec{\sigma}$, we have $\text{PF}_m(f') \leq \text{PF}_m(f)$, as required. \square

We are now equipped to prove a lower bound on proportional fairness of any strategyproof voting rule.

THEOREM D.8. *For every strategyproof voting rule f , we have $\text{PF}_m(f) = \Omega(\sqrt{m})$.*

PROOF. Let f be a strategyproof voting rule. Due to Lemma D.7, we may assume that f is anonymous and neutral. Due to Proposition D.6, f is a probability mixture that implements some point-voting scheme characterized by \vec{w} with probability $p \in [0, 1]$ and some supporting-size scheme characterized by \vec{z} with probability $1 - p$.

Let $\alpha = \text{PF}_m(f)$. We will show that $\alpha = \Omega(\sqrt{m})$. Similarly to Bhaskar et al. [14], we construct a sequence of profiles $\vec{\sigma}^r$, one for each $r \in [m]$, and show that f must select a distribution that is at least $\Omega(\sqrt{m})$ -proportionally fair on at least one of these profiles, so that $\alpha = \Omega(\sqrt{m})$.

Fix any $r \in [m]$ and construct the preference profile $\vec{\sigma}^r$ with $n = m - 1$ votes,¹⁵ such that a special alternative a^* appears in position r in all the votes, while the remaining $m - 1$ alternatives fill the remaining positions in a cyclic order. Let $\mathbf{x}^r = f(\vec{\sigma}^r)$. Then,

$$\begin{aligned} \mathbf{x}^r(a^*) &= p \cdot \frac{1}{n} \cdot \sum_{i \in N} w_{\sigma_i(a^*)} + (1 - p) \cdot \frac{1}{\binom{m}{2}} \cdot \sum_{a \in A \setminus \{a^*\}} z_{n_{a^* > a}} \\ &\leq 1 \cdot \frac{1}{n} \cdot \sum_{i \in N} w_r + 1 \cdot \frac{1}{\binom{m}{2}} \cdot \sum_{a \in A \setminus \{a^*\}} 1 \\ &= w_r + \frac{2}{m}. \end{aligned} \tag{8}$$

¹⁵The profile can easily be replicated to make n any multiple of $m - 1$.

As argued by Bhaskar et al. [14], this observation would effectively allow us to ignore the impact of the supporting-size scheme, which can only select the “desired” alternative a^* with a small probability of $2/m$, and focus on the impact of the point-voting scheme. Let us lower bound proportional fairness of \mathbf{x}^r on $\vec{\sigma}^r$. Now, we have

$$\begin{aligned}
\text{PF}(\mathbf{x}^r, \vec{\sigma}^r) &= \max_{a \in A} \frac{1}{n} \sum_{i \in N} \frac{1}{\mathbf{x}^r(h_i(a))} && \text{(by Lemma 4.1)} \\
&\geq \frac{1}{n} \sum_{i \in N} \frac{1}{\mathbf{x}^r(h_i(a^*))} \\
&\geq \frac{n}{\sum_{i \in N} \mathbf{x}^r(h_i(a^*))} && \text{(by the AM-HM inequality)} \\
&= \frac{n}{n \cdot (\mathbf{x}^r(a^*) + \frac{r-1}{m-1} \cdot (1 - \mathbf{x}^r(a^*)))} \\
&\geq \frac{1}{\mathbf{x}^r(a^*) + \frac{r}{m} \cdot (1 - \mathbf{x}^r(a^*))} \\
&\geq \frac{1}{\mathbf{x}^r(a^*) + \frac{r}{m}},
\end{aligned}$$

where the fourth transition holds because anonymity and neutrality of f implies $\mathbf{x}^r(a) = \frac{1 - \mathbf{x}^r(a^*)}{m-1}$ for each $a \in A \setminus \{a^*\}$. By definition of α , we have $\alpha \geq \text{PF}(\mathbf{x}^r, \vec{\sigma}^r)$. Thus it follows that

$$\mathbf{x}^r(a^*) + \frac{r}{m} \geq \frac{1}{\alpha}.$$

Using Equation (8), we have

$$w_r \geq \max\left(\frac{1}{\alpha} - \frac{r+2}{m}, 0\right).$$

Summing over all $r \in [m]$, we get

$$\begin{aligned}
1 &= \sum_{r \in [m]} w_r \geq \sum_{r \in [m]} \max\left(\frac{1}{\alpha} - \frac{r+2}{m}, 0\right) \\
&= \sum_{r=1}^{\lfloor \frac{m}{\alpha} \rfloor - 2} \left(\frac{1}{\alpha} - \frac{r+2}{m}\right) \\
&\geq \frac{\frac{m}{\alpha} - 3}{\alpha} - \frac{\frac{m}{\alpha} \left(\frac{m}{\alpha} + 1\right) - 6}{2m} \\
&\geq \frac{m}{2\alpha^2} - \frac{7}{2\alpha} + \frac{3}{m},
\end{aligned}$$

where the fourth transition uses $m/\alpha - 1 \leq \lfloor m/\alpha \rfloor \leq m/\alpha$. The above expression simplifies to

$$(2m - 6)\alpha^2 + 7m\alpha - m^2 \geq 0,$$

which shows

$$\alpha \geq \frac{-7m + \sqrt{49m^2 + 4m^2(2m - 6)}}{2(2m - 6)} = \Omega(\sqrt{m}),$$

as needed. □

E PROPORTIONAL FAIRNESS VIA APPROXIMATELY STABLE COMMITTEES

As indicated in Section 4, we prove that a rule similar to our stable lottery rule (f_{SLR}), which uses a deterministic committee satisfying approximate stability instead of a lottery over committees satisfying exact stability, is $O(\sqrt{m})$ -proportionally fair, and therefore achieves $O(\sqrt{m})$ distortion with respect to the Nash welfare. Let us first formally introduce approximate stability for committees.

Definition E.1 (Approximately Stable Committees). For a committee X with $|X| = k$ and an alternative a^* , recall that $V(a^*, X) = |\{i \in N : a^* \succ_i X\}|$ denotes the number of agents who prefer a^* to all alternatives in X . We say that X is c -stable if for all alternatives $a^* \notin X$, we have $V(a^*, X) \leq c \cdot \frac{n}{k}$.

Note that 1-stable committees are precisely the stable committees introduced in Section 3.1. As mentioned in that section, there exist preference profiles and sizes k where no stable committee of size k exists [36, Thm. 4]. However, by derandomizing the stable lottery of Cheng et al. [26], Jiang et al. [36] proved the following:

THEOREM E.2 (JIANG ET AL. 36). *Given any ranked preference profile and $k \in [m]$, a 16-stable committee of size k exists and a $(16 + \varepsilon)$ -stable committee of size k can be computed in $\text{poly}(n, m, 1/\varepsilon)$ time for sufficiently small constant $\varepsilon > 0$.*

Let us introduce a voting rule that uses an approximately stable committee in the same manner in which the rule f_{SLR} from Section 3 uses an exactly stable lottery. Note that despite the use of a deterministic committee, the rule is still probabilistic in the end.

Definition E.3 (c -Stable Committee Rule, $f_{c\text{-SCR}}$). Let X be a c -stable committee of size $k = \sqrt{m}$. The c -Stable Lottery Rule ($f_{c\text{-SCR}}$) works as follows: With probability $1/2$, choose an alternative uniformly at random from X , and with probability $1/2$, choose an alternative uniformly at random from A . Therefore, each alternative $a \in A$ is selected with probability $\mathbf{x}(a) = \frac{1}{2\sqrt{m}} \cdot \mathbb{I}[a \in X] + \frac{1}{2m}$, where \mathbb{I} is the indicator function.

Next, we prove that $f_{c\text{-SCR}}$ is $O(\sqrt{m})$ -proportionally fair when c is constant.

THEOREM E.4. *For constant c , we have $\text{PF}_m(f_{c\text{-SCR}}) = O(\sqrt{m})$.*

PROOF. For constant c , consider the $f_{c\text{-SCR}}$ rule. Fix an arbitrary preference profile $\vec{\sigma}$. Let X be the c -stable committee of size \sqrt{m} that our rule uses to output distribution \mathbf{x} on this profile. We want to prove that $\text{PF}(\mathbf{x}, \vec{\sigma}) = O(\sqrt{m})$.

By Lemma 4.1, we have that

$$\text{PF}(\mathbf{x}, \vec{\sigma}) = \max_{a \in A} \frac{1}{n} \sum_{i \in N} \frac{1}{\mathbf{x}(h_i(a))} = \frac{1}{n} \sum_{i \in N} \frac{1}{\mathbf{x}(h_i(a^*))}, \quad (9)$$

where $\mathbf{x}(h_i(a))$ is the probability placed on the set of alternatives $h_i(a) = \{a' : a' \succsim_i a\}$ under \mathbf{x} and a^* is taken to be an arg max.

Let $S \subseteq N$ denote the set of $V(a^*, X)$ many agents who prefer a^* to every alternative in X . Because X is c -stable, we know that $|S| = V(a^*, X) \leq c \cdot n/\sqrt{m}$. By definition of S , each agent $i \in N \setminus S$ satisfies $h_i(a^*) \cap X \neq \emptyset$, implying that $\mathbf{x}(h_i(a^*)) \geq \frac{1}{2\sqrt{m}}$. For each agent $i \in S$, we have

$\mathbf{x}(h_i(a^*)) \geq \mathbf{x}(a^*) \geq \frac{1}{2m}$. Plugging these in Equation (9), we get

$$\begin{aligned}
\text{PF}(\mathbf{x}, \vec{\sigma}) &= \frac{1}{n} \sum_{i \in N} \frac{1}{\mathbf{x}(h_i(a^*))} \\
&= \frac{1}{n} \left(\sum_{i \in S} \frac{1}{\mathbf{x}(h_i(a^*))} + \sum_{i \in N \setminus S} \frac{1}{\mathbf{x}(h_i(a^*))} \right) \\
&\leq \frac{1}{n} \left(|S| \cdot 2m + |N \setminus S| \cdot 2\sqrt{m} \right) \\
&\leq \frac{1}{n} \left(c \cdot \frac{n}{\sqrt{m}} \cdot 2m + n \cdot 2\sqrt{m} \right) \\
&= 2 \cdot (c + 1) \cdot \sqrt{m} = O(\sqrt{m}).
\end{aligned}$$

This completes the proof. \square

Note that an upper bound on proportional fairness also applies to distortion with respect to the Nash welfare.

COROLLARY E.5. *For constant c , the distortion of $f_{c\text{-SCR}}$ with respect to the Nash welfare is $D^{\text{NW}}(f_{c\text{-SCR}}, \mathcal{U}^{\text{all}}) = O(\sqrt{m})$.*

For distortion with respect to the utilitarian social welfare, the proof of Theorem 3.4 can easily be modified to show that for each constant c , we have $D^{\text{UW}}(f_{c\text{-SCR}}, \mathcal{U}^{\text{balanced}}) = O(\sqrt{m})$, i.e., the distortion of $f_{c\text{-SCR}}$ is $O(\sqrt{m})$ for balanced utilities (and therefore also for unit-sum and unit-range utilities). In the main body, we focused on the Stable Lottery Rule instead of the Stable Committee Rule because for the former it is easier to prove existence, its output is easier to compute, and due to the existence of an *exact* stable lottery (as opposed to an approximately stable committee), we get a distortion upper bound of $2\sqrt{m}$, which is close to the lower bound of $\sqrt{m}/2$.¹⁶

We end this section by recalling that in Section 4, we show that a better performance on proportional fairness can be achieved, using a different technique based on the minimax theorem.

F LOWER BOUND ON DISTORTION WITH RESPECT TO NASH WELFARE

In this section, we provide the proof of Proposition 4.7, which provides a lower bound on the distortion of any voting rule with respect to Nash welfare.

PROPOSITION. *Every voting rule has distortion at least $(\frac{m^m}{m!})^{\frac{1}{m}}$ with respect to the Nash welfare.*

PROOF. Consider the cyclic profile $\vec{\sigma}$ from the proof of Theorem 4.6. Fix any distribution \mathbf{x} and let $T_i(\ell)$ be defined as in the proof of Theorem 4.6. Then, using the same argument as before, we

¹⁶Boutilier et al. [17] prove a lower bound of $\sqrt{m}/3$, but a careful look at their analysis shows that it actually yields a lower bound of $\sqrt{m}/2$.

have

$$\begin{aligned}
D^{\text{NW}}(\mathbf{x}, \vec{\sigma}) &= \max_{\mathbf{y} \in \Delta(A)} \left(\prod_{i \in [m]} \sup_{u_i: u_i \succ \sigma_i} \frac{u_i(\mathbf{y})}{u_i(\mathbf{x})} \right)^{1/m} \geq \max_{a \in A} \left(\prod_{i \in [m]} \sup_{u_i: u_i \succ \sigma_i} \frac{u_i(a)}{u_i(\mathbf{x})} \right)^{1/m} \\
&= \max_{a \in A} \left(\prod_{i \in [m]} \frac{1}{\mathbf{x}(h_i(a))} \right)^{1/m} \geq \left(\prod_{a \in A} \prod_{i \in [m]} \frac{1}{\mathbf{x}(h_i(a))} \right)^{1/m^2} \\
&= \left(\prod_{\ell \in [m]} \prod_{i \in [m]} \frac{1}{T_i(\ell)} \right)^{1/m^2} \geq \left(\prod_{\ell \in [m]} \frac{m}{\sum_{i \in [m]} T_i(\ell)} \right)^{1/m} = \left(\prod_{\ell \in [m]} \frac{m}{\ell} \right)^{1/m},
\end{aligned}$$

where, crucially, the penultimate transition uses the standard inequality between geometric and harmonic means. The final expression is the desired lower bound. \square

G THE CASE OF TWO ALTERNATIVES

In this section, we analyze the case of exactly two alternatives, say, a_1 and a_2 . For each of the objectives that we studied in this paper (distortion with unit-sum utilities, distortion with unit-range utilities, distortion with respect to Nash welfare, proportional fairness), we will explicitly write down a voting rule that, for every preference profile, selects the instance-optimal distribution. We will also compute the worst-case distortion/proportional fairness of these rules and hence the best possible values obtainable in the two-alternative case. (These values turn out to be $3/2$, $4/3$, $\sqrt{2}$, and $3/2$, respectively.)

As there are only two possible rankings over two alternatives, we can summarize a preference profile $\vec{\sigma}$ by a real number $\alpha(\vec{\sigma}) \in [0, 1]$, which denotes the fraction of agents who prefer a_1 to a_2 ; then, the remaining $1 - \alpha(\vec{\sigma})$ fraction of agents prefer a_2 to a_1 . Similarly, the outcome of a voting rule on a preference profile $\vec{\sigma}$ can also be viewed as a real number $\beta(\vec{\sigma}) \in [0, 1]$ which is the probability placed on a_1 . We will adapt our notation accordingly throughout this section.

We can disregard the extreme cases where $\alpha(\vec{\sigma}) \in \{0, 1\}$ because in such cases, one alternative is the more preferred one for all the agents (a_1 if $\alpha(\vec{\sigma}) = 1$ and a_2 if $\alpha(\vec{\sigma}) = 0$); the voting rule should clearly choose it as this achieves the optimal distortion of 1 with respect to any welfare function for the class \mathcal{U}^{all} of all utility functions as well as the optimal 1-proportional fairness.

G.1 Distortion With Unit-Sum Utilities

Given a preference profile $\vec{\sigma}$, the voting rule $f_{2\text{-UW}}^{\text{unit-sum}}$ is defined to choose $\beta(\vec{\sigma})$ satisfying

$$\frac{\beta(\vec{\sigma})}{(1 - \beta(\vec{\sigma}))} = \frac{(2 - \alpha(\vec{\sigma})) \cdot \alpha(\vec{\sigma})}{(1 - \alpha(\vec{\sigma})) \cdot (1 + \alpha(\vec{\sigma}))}. \quad (10)$$

In other words, the rule sets $\beta(\vec{\sigma}) = \frac{(2 - \alpha(\vec{\sigma})) \cdot \alpha(\vec{\sigma})}{1 + 2\alpha(\vec{\sigma}) \cdot (1 - \alpha(\vec{\sigma}))}$.

THEOREM G.1. *For $m = 2$ alternatives, and for any preference profile $\vec{\sigma}$, the voting rule $f_{2\text{-UW}}^{\text{unit-sum}}$ selects a distribution \mathbf{x} minimizing $D^{\text{UW}}(\mathbf{x}, \vec{\sigma}, \mathcal{U}^{\text{unit-sum}})$.*

The rule achieves distortion $D_2^{\text{UW}}(f_{2\text{-UW}}^{\text{unit-sum}}, \mathcal{U}^{\text{unit-sum}}) = 3/2$, which is the best possible among all voting rules.

PROOF. Fix a preference profile $\vec{\sigma}$. Let us write $\alpha = \alpha(\vec{\sigma})$ and let

$$D^{\text{UW}}(\alpha) = \min_{\mathbf{x} \in \Delta(A)} D^{\text{UW}}(\mathbf{x}, \vec{\sigma}, \mathcal{U}^{\text{unit-sum}})$$

be the best achievable distortion on profile $\vec{\sigma}$ with unit-sum utilities. Then,

$$\begin{aligned} D^{\text{UW}}(\alpha) &= \min_{\beta \in [0,1]} \max_{\vec{u} \in (\mathcal{U}^{\text{unit-sum}})^n: \vec{u} \succ \vec{\sigma}} \frac{\max\{\text{UW}(a_1, \vec{u}), \text{UW}(a_2, \vec{u})\}}{\beta \cdot \text{UW}(a_1, \vec{u}) + (1 - \beta) \cdot \text{UW}(a_2, \vec{u})} \\ &= \min_{\beta \in [0,1]} \max \left\{ \max_{\vec{u} \in (\mathcal{U}^{\text{unit-sum}})^n: \vec{u} \succ \vec{\sigma}} \frac{\text{UW}(a_1, \vec{u})}{\beta \cdot \text{UW}(a_1, \vec{u}) + (1 - \beta) \cdot \text{UW}(a_2, \vec{u})}, \right. \\ &\quad \left. \max_{\vec{u} \in (\mathcal{U}^{\text{unit-sum}})^n: \vec{u} \succ \vec{\sigma}} \frac{\text{UW}(a_2, \vec{u})}{\beta \cdot \text{UW}(a_1, \vec{u}) + (1 - \beta) \cdot \text{UW}(a_2, \vec{u})} \right\}. \end{aligned} \quad (11)$$

Worst-case utilities. In Equation (11), let us analyze the worst-case utility profile \vec{u} for the term with the numerator $\text{UW}(a_1, \vec{u})$. Using the fact that $\text{UW}(a_1, \vec{u}) + \text{UW}(a_2, \vec{u}) = n$ (by the unit-sum assumption), the expression of interest in Equation (11) is

$$\frac{\text{UW}(a_1, \vec{u})}{\beta \cdot \text{UW}(a_1, \vec{u}) + (1 - \beta)(n - \text{UW}(a_1, \vec{u}))} = \frac{1}{\beta + (1 - \beta)(n - \text{UW}(a_1, \vec{u}))/\text{UW}(a_1, \vec{u})},$$

which is maximized when $\text{UW}(a_1, \vec{u})$ is maximized. To obtain the highest utilitarian welfare for a_1 , α fraction of agents i who rank a_1 first must have $(u_i(a_1), u_i(a_2)) = (1, 0)$ and the remaining $1 - \alpha$ fraction of agents i who rank a_1 second must have $(u_i(a_2), u_i(a_1)) = (1/2, 1/2)$. Similarly, the worst-case utility profile \vec{u} for the term with the numerator $\text{UW}(a_2, \vec{u})$ is achieved when the α fraction of agents i who rank a_2 second have $(u_i(a_1), u_i(a_2)) = (1/2, 1/2)$ and the remaining $1 - \alpha$ fraction of agents i who rank a_2 first have $(u_i(a_2), u_i(a_1)) = (1, 0)$. Therefore, Equation (11) becomes

$$\begin{aligned} D^{\text{UW}}(\alpha) &= \min_{\beta \in [0,1]} \max \left\{ \frac{\alpha + \frac{1-\alpha}{2}}{(\alpha + \frac{1-\alpha}{2})\beta + \frac{1-\alpha}{2} \cdot (1-\beta)}, \frac{\frac{\alpha}{2} + 1 - \alpha}{(\frac{\alpha}{2} + 1 - \alpha)(1-\beta) + \frac{\alpha}{2} \cdot \beta} \right\} \\ &= \min_{\beta \in [0,1]} \max \left\{ \frac{1}{\beta + \frac{(1-\alpha)/2}{\alpha + (1-\alpha)/2} \cdot (1-\beta)}, \frac{1}{\frac{\alpha/2}{\alpha/2 + 1 - \alpha} \cdot \beta + 1 - \beta} \right\}. \end{aligned}$$

The left term in the maximum is decreasing in β while the right term is increasing. Hence, the maximum of the two is minimized when the two terms become equal, i.e.,

$$\begin{aligned} \beta + \frac{(1-\alpha)/2}{\alpha + (1-\alpha)/2} \cdot (1-\beta) &= \frac{\alpha/2}{\alpha/2 + 1 - \alpha} \cdot \beta + 1 - \beta \\ \iff \beta \cdot \left(\frac{1}{\frac{\alpha/2}{1-\alpha} + 1} \right) &= (1-\beta) \cdot \left(\frac{1}{\frac{1-\alpha}{2\alpha} + 1} \right) \iff \frac{1-\beta}{\beta} = \frac{(1-\alpha)(1+\alpha)}{(2-\alpha)\alpha}, \end{aligned}$$

which precisely yields the β satisfying Equation (10), proving that $f_{2-\text{UW}}^{\text{unit-sum}}$ returns the distribution with the best possible distortion on profile $\vec{\sigma}$. Further, at this value of β , the distortion achieved is $D_2^{\text{UW}}(f_{2-\text{UW}}^{\text{unit-sum}}) = \max_{\alpha \in [0,1]} D^{\text{UW}}(\alpha) = 3/2$, which is attained at $\alpha = 1/2$ (for which the optimal β is also $\beta = 1/2$). \square

From Section 3, we know that, with m alternatives, the optimal distortion with respect to utilitarian welfare and unit-sum utilities is $\Theta(\sqrt{m})$, where the constant hidden in the asymptotic notation lies in $[\frac{1}{2}, 2]$. Interestingly, Theorem G.1 shows that the optimal distortion for $m = 2$ is $\frac{3}{2} = \frac{3}{2\sqrt{2}} \cdot \sqrt{m}$, leading one to wonder whether the constant $\frac{3}{2\sqrt{2}} \approx 1.06066$ is (close to) the true constant for this problem.

G.2 Distortion With Unit-Range Utilities

Given a preference profile $\vec{\sigma}$, the voting rule $f_{2\text{-UW}}^{\text{unit-range}}$ is defined to choose $\beta(\vec{\sigma}) = \alpha(\vec{\sigma})$.

THEOREM G.2. *For $m = 2$ alternatives, and for any preference profile $\vec{\sigma}$, the voting rule $f_{2\text{-UW}}^{\text{unit-range}}$ selects a distribution \mathbf{x} minimizing $D^{\text{UW}}(\mathbf{x}, \vec{\sigma}, \mathcal{U}^{\text{unit-range}})$.*

The rule achieves distortion $D_2^{\text{UW}}(f_{2\text{-UW}}^{\text{unit-range}}, \mathcal{U}^{\text{unit-range}}) = 4/3$, which is the best possible among all voting rules.

PROOF. Fix a preference profile $\vec{\sigma}$. Let us write $\alpha = \alpha(\vec{\sigma})$ and let

$$D^{\text{UW}}(\alpha) = \min_{\mathbf{x} \in \Delta(A)} D^{\text{UW}}(\mathbf{x}, \vec{\sigma}, \mathcal{U}^{\text{unit-range}})$$

be the best achievable distortion on profile $\vec{\sigma}$ with unit-range utilities. Then, as in the proof of Theorem G.1, we have

$$\begin{aligned} D^{\text{UW}}(\alpha) &= \min_{\beta \in [0,1]} \max_{\vec{u} \in (\mathcal{U}^{\text{unit-range}})^n: \vec{u} \succ \vec{\sigma}} \frac{\max\{\text{UW}(a_1, \vec{u}), \text{UW}(a_2, \vec{u})\}}{\beta \cdot \text{UW}(a_1, \vec{u}) + (1 - \beta) \cdot \text{UW}(a_2, \vec{u})} \\ &= \min_{\beta \in [0,1]} \max \left\{ \begin{aligned} &\max_{\vec{u} \in (\mathcal{U}^{\text{unit-range}})^n: \vec{u} \succ \vec{\sigma}} \frac{\text{UW}(a_1, \vec{u})}{\beta \cdot \text{UW}(a_1, \vec{u}) + (1 - \beta) \cdot \text{UW}(a_2, \vec{u})}, \\ &\max_{\vec{u} \in (\mathcal{U}^{\text{unit-range}})^n: \vec{u} \succ \vec{\sigma}} \frac{\text{UW}(a_2, \vec{u})}{\beta \cdot \text{UW}(a_1, \vec{u}) + (1 - \beta) \cdot \text{UW}(a_2, \vec{u})} \end{aligned} \right\}. \quad (12) \end{aligned}$$

Worst-case utilities. The only difference compared to the proof of Theorem G.1 is the analysis of the worst-case utility profiles in the two expressions inside the maximum in Equation (12). By Lemma B.1, we know that the worst-case utility profile is an approval utility profile. All agents approve their first-ranked alternative and the only question is whether they also approve their second-ranked alternative.

Consider the first term inside the maximum in Equation (12) with $\text{UW}(a_1, \vec{u})$ in the numerator. We want to find the utility profile \vec{u} that maximizes this term. For the α fraction of agents who rank a_1 above a_2 , we can assume wlog that $u_i(a_2) = 0$ as it can only increase this term. For the rest of the agents i that rank a_2 above a_1 , we know $u_i(a_2) = 1$. Hence, $\text{UW}(a_2, \vec{u}) = (1 - \alpha) \cdot n$. Given this fixed value of $\text{UW}(a_2, \vec{u})$, note that $\frac{\text{UW}(a_1, \vec{u})}{\beta \cdot \text{UW}(a_1, \vec{u}) + (1 - \beta) \cdot \text{UW}(a_2, \vec{u})} = \frac{1}{\beta + (1 - \beta) \cdot \frac{\text{UW}(a_2, \vec{u})}{\text{UW}(a_1, \vec{u})}}$ is increasing in $\text{UW}(a_1, \vec{u})$. Hence, the term is maximized when $\text{UW}(a_1, \vec{u})$ is the highest, meaning that the agents who rank a_2 above a_1 also approve a_1 .

Similarly, for the term in Equation (12) with $\text{UW}(a_2, \vec{u})$ in the numerator, the worst-case utility profile is as follows: α fraction of agents i who rank a_1 first have $(u_i(a_1), u_i(a_2)) = (1, 1)$ and the remaining $1 - \alpha$ fraction of agents i who rank a_2 first have $(u_i(a_2), u_i(a_1)) = (1, 0)$. Therefore, Equation (12) becomes

$$D^{\text{UW}}(\alpha) = \min_{\beta \in [0,1]} \max \left\{ \frac{1}{\beta + (1 - \alpha)(1 - \beta)}, \frac{1}{\alpha\beta + (1 - \beta)} \right\}$$

The left term in the maximum is decreasing in β while the right term is increasing. Hence, the maximum of the two terms is minimized when the two terms are equal, which precisely yields the $\beta = \alpha$. This proves that $f_{2\text{-UW}}^{\text{unit-range}}$ returns the distribution with the best possible distortion on profile $\vec{\sigma}$. Further, at this value of β , the distortion achieved is $D_2^{\text{UW}}(f_{2\text{-UW}}^{\text{unit-range}}) = \max_{\alpha \in [0,1]} D^{\text{UW}}(\alpha) = 4/3$, which is attained at $\alpha = \beta = 1/2$. \square

G.3 Distortion With Respect to Nash Welfare

Given a preference profile $\vec{\sigma}$, the voting rule $f_{2\text{Nash}}$ is defined to choose $\beta(\vec{\sigma})$ satisfying

$$\frac{\ln(\beta(\vec{\sigma}))}{\ln(1 - \beta(\vec{\sigma}))} = \frac{1 - \alpha(\vec{\sigma})}{\alpha(\vec{\sigma})}. \quad (13)$$

In other words, the rule sets $\beta(\vec{\sigma}) = g^{-1}(\alpha(\vec{\sigma}))$, where $g(x) = \frac{\ln(1-x)}{\ln(x) + \ln(1-x)}$.

THEOREM G.3. *For $m = 2$ alternatives, and for any preference profile $\vec{\sigma}$, the voting rule $f_{2\text{Nash}}$ selects a distribution \mathbf{x} minimizing $D^{\text{NW}}(\mathbf{x}, \vec{\sigma}, \mathcal{U}^{\text{all}})$, the distortion with respect to Nash welfare.*

The rule achieves distortion $D_2^{\text{NW}}(f_{2\text{Nash}}, \mathcal{U}^{\text{all}}) = \sqrt{2}$ with respect to Nash welfare, which is the best possible among all voting rules.

PROOF. The lower bound $\sqrt{2}$ for all voting rules was proved in Proposition 4.7 (see Appendix F). Hence, it remains to prove the upper bound for $f_{2\text{Nash}}$. Fix a preference profile $\vec{\sigma}$. Let us write $\alpha = \alpha(\vec{\sigma})$ and let $D^{\text{NW}}(\alpha) = \min_{\mathbf{x} \in \Delta(A)} D^{\text{NW}}(\mathbf{x}, \vec{\sigma}, \mathcal{U}^{\text{all}})$ be the best achievable distortion on profile $\vec{\sigma}$. Note that

$$D^{\text{NW}}(\alpha) = \min_{\mathbf{x} \in \Delta(A)} \max_{\mathbf{y} \in \Delta(A)} \sup_{\vec{u} \in (\mathcal{U}^{\text{all}})^n: \vec{u} \succ \vec{\sigma}} \left(\prod_{i \in N} \frac{u_i(\mathbf{y})}{u_i(\mathbf{x})} \right)^{1/n}. \quad (14)$$

Reduction to approval utilities. From Lemma B.3, we know that the worst case for distortion with respect to Nash welfare is achieved at an approval utility profile. Under an approval utility profile \vec{u} , the α fraction of agents i who prefer a_1 to a_2 have $(u_i(a_1), u_i(a_2))$ equal to $(1, 0)$ or $(1, 1)$, and the remaining $1 - \alpha$ fraction of agents i have $(u_i(a_1), u_i(a_2))$ equal to $(0, 1)$ or $(1, 1)$. If an agent approves both alternatives, then $\frac{u_i(\mathbf{y})}{u_i(\mathbf{x})} = 1$ regardless of \mathbf{x} and \mathbf{y} . Based on these observations, we can rewrite $D^{\text{NW}}(\alpha)$ from Equation (14) as

$$D^{\text{NW}}(\alpha) = \min_{\beta \in [0,1]} \max_{z \in [0,1]} \left(\max \left\{ \frac{z}{\beta}, 1 \right\}^\alpha \cdot \max \left\{ \frac{1-z}{1-\beta}, 1 \right\}^{1-\alpha} \right).$$

Finding the optimal distribution. If $z > \beta$, then $\frac{1-z}{1-\beta} < 1$ and the inner expression evaluates to $(\frac{z}{\beta})^\alpha$, which is maximized at $z = 1$. Similarly, when $z < \beta$, the inner expression evaluates to $(\frac{1-z}{1-\beta})^{1-\alpha}$, which is maximized at $z = 0$. Therefore, we have

$$D^{\text{NW}}(\alpha) = \min_{\beta \in [0,1]} \max \left\{ \left(\frac{1}{\beta} \right)^\alpha, \left(\frac{1}{1-\beta} \right)^{1-\alpha} \right\}.$$

One can check that the unique minimizer of this expression is precisely the β satisfying Equation (13), which our rule $f_{2\text{Nash}}$ chooses. Further, at this β , the distortion achieved is $D_2^{\text{NW}}(f_{2\text{Nash}}) = \max_{\alpha \in [0,1]} D^{\text{NW}}(\alpha) = \sqrt{2}$, which is attained at $\alpha = 1/2$ (for which the optimal β is also $\beta = 1/2$). \square

G.4 Proportional Fairness

Given a preference profile $\vec{\sigma}$, the voting rule $f_{2\text{PF}}$ is defined to choose $\beta(\vec{\sigma})$ satisfying

$$\frac{\beta(\vec{\sigma})^2}{(1 - \beta(\vec{\sigma}))^2} = \frac{\alpha(\vec{\sigma})}{1 - \alpha(\vec{\sigma})}. \quad (15)$$

In other words, the rule sets $\beta(\vec{\sigma}) = \frac{\sqrt{\alpha(\vec{\sigma})}}{\sqrt{1-\alpha(\vec{\sigma})} + \sqrt{\alpha(\vec{\sigma})}}$.

THEOREM G.4. For $m = 2$ alternatives, and for any preference profile $\vec{\sigma}$, the voting rule $f_{2\text{PF}}$ selects a distribution \mathbf{x} minimizing $\text{PF}(\mathbf{x}, \vec{\sigma})$.

The rule is $(3/2)$ -proportionally fair, which is the best possible among all voting rules.

PROOF. The lower bound for all voting rules proved in Theorem 4.6 is the desired lower bound of $H_2 = 3/2$ when $m = 2$. Hence, it remains to prove the upper bound for $f_{2\text{PF}}$. Fix a preference profile $\vec{\sigma}$. Let us write $\alpha = \alpha(\vec{\sigma})$ and let $\text{PF}(\alpha) = \min_{\mathbf{x} \in \Delta(A)} \text{PF}(\mathbf{x}, \vec{\sigma})$. From Lemma 4.1, we have

$$\begin{aligned} \text{PF}(\alpha) &= \min_{\mathbf{x} \in \Delta(A)} \max_{a \in A} \frac{1}{n} \sum_{i \in N} \frac{1}{\mathbf{x}(h_i(a))} \\ &= \min_{\beta \in [0,1]} \max \left\{ \frac{\alpha}{\beta} + (1 - \alpha), \alpha + \frac{1 - \alpha}{1 - \beta} \right\}. \end{aligned}$$

This expression is minimized when

$$\frac{\alpha}{\beta} + (1 - \alpha) = \alpha + \frac{1 - \alpha}{1 - \beta} \iff \alpha \cdot \frac{1 - \beta}{\beta} = (1 - \alpha) \cdot \frac{\beta}{1 - \beta} \iff \frac{\alpha}{1 - \alpha} = \frac{\beta^2}{(1 - \beta)^2},$$

which precisely yields the β satisfying Equation (15), proving that $f_{2\text{PF}}$ returns the distribution with the best possible proportional fairness on profile $\vec{\sigma}$. Further, at this value of β , the proportional fairness achieved is $\text{PF}_2(f_{2\text{PF}}) = \max_{\alpha \in [0,1]} \text{PF}(\alpha) = 3/2$, which is attained at $\alpha = 1/2$ (for which the optimal β is also $\beta = 1/2$). \square