

Optimized Distortion and Proportional Fairness in Voting

SOROUGH EBADIAN, University of Toronto

ANSON KAHNG, University of Toronto

DOMINIK PETERS, University of Toronto

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 produced by the outcome selected by the voting rule to the social welfare of the optimum outcome, for the worst possible input. The previous literature has studied distortion with (normalized) unit-sum utility functions, 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. We show that this bound holds even for a larger class of utilities, including approval (0/1) utilities, and for distortion with respect to another common social welfare function, namely the Nash welfare.

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 proportional fairness $O(\log m)$, where lower is better and 1 is optimal. This voting rule thereby achieves $O(\log m)$ distortion with respect to the Nash welfare, and selects a distribution that is approximately stable by being in the $O(\log m)$ -approximate core, making it interesting for applications in participatory budgeting.

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 .

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 particular 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]. We propose the first voting rule achieving the asymptotically optimal distortion $O(\sqrt{m})$, matching the lower bound by Boutilier et al. [2015] and resolving an important open question in this line of work. Our proof shows that the same rule is also optimal for the class of *approval utilities*, in which each agent has utility 1 for a subset of alternatives and utility 0 for the rest. Such utilities are commonly studied in voting (though, to the best of our knowledge, not in the context of distortion) because approval utilities naturally correspond to the commonly used approval voting method. The rule can be computed in polynomial time, but we also show that we can compute an instance-optimal distribution in polynomial time for the utility classes that we study.

Interestingly, while our new voting rule for distortion achieves high social welfare, it internally aims for a fair outcome. Specifically, it uses tools from multi-winner voting designed for selecting a *representative* committee (i.e., subset of alternatives), in which, informally, as many agents as possible 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 some 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. In particular, 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 agent or any 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 with proportional fairness 1 (lower is better). We show that given only a preference profile of rankings, there always exists a distribution that achieves proportional fairness $O(\log m)$ regardless of the underlying utility functions (consistent with the input rankings). We show that our result is optimal, because there are preference profiles for which every distribution has proportional fairness $\Omega(\log m)$ under some consistent utility functions. Our proof for the upper bound uses the minimax theorem from the theory of zero-sum games, and we show that an instance-optimal distribution can be computed in polynomial time.

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. 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; they wish to see more money spent on higher-ranked projects. 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 α -core. Our rule for proportional fairness selects an outcome that is in the $O(\log m)$ -approximate core.

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 analyzed the distortion of many common voting rules. Caragiannis and Procaccia [2011] also evaluated the distortion of prominent voting rules, but from the perspective of optimizing embeddings, which (perhaps randomly) map cardinal utilities to ordinal votes that voting rules take as input. Later work designed new voting rules with low distortion for both single-winner [Boutilier et al., 2015] and multi-winner settings [Caragiannis et al., 2017]. In particular, Boutilier et al. [2015] derive an upper bound of $O(\sqrt{m} \log^* m)$ and a lower bound of $\Omega(\sqrt{m})$ on the optimal distortion for probabilistic voting rules.¹ Our $O(\sqrt{m})$ upper bound matches their lower bound and eliminates the $\log^* m$ gap.

Going beyond single-winner voting, Caragiannis et al. [2017] studied 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 presumed that the utility of an agent for a committee is the maximum utility of the agent for any alternative in the committee. They proved that the optimal distortion of deterministic rules is $\Theta(1 + m(m - k)/k)$ (implying an optimal distortion bound of $\Theta(m^2)$ for single-winner voting), but left a gap of $\Theta(m^{1/6})$ between their upper and lower bounds for the optimal distortion of probabilistic rules.

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 tradeoff 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] characterized the Pareto frontier of this tradeoff, showing that in order to achieve distortion d , probabilistic voting rules require voters to communicate only $\tilde{\Theta}(m/d^3)$ bits of information whereas deterministic voting rules require $\tilde{\Theta}(m/d)$ bits, establishing probabilistic rules as superior in this context. Amanatidis

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

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 proved that asking only $O(\log^2 m)$ value queries or $O(\log^2 m)$ comparison queries is sufficient to achieve constant distortion. Benade et al. [2021] studied participatory budgeting, in which each alternative has a cost, the goal is to find a subset of alternatives with total cost at most a given budget, and the utility of an agent for a set of alternatives is the sum of her utilities for the alternatives in the set. They compared four protocols for eliciting agent preferences, and proved that threshold approval votes (which ask agents to identify alternatives for which their utility is at least a threshold) lead to the lowest distortion of $O(\log^2 m)$ (as compared to $O(\sqrt{m} \cdot \log m)$ for ranked preferences). Filos-Ratsikas et al. [2020] studied distortion in settings where agents are partitioned into districts and vote over local alternatives, after which the overall winner is chosen from local winners. They proved that such distributed elections lead to higher distortion.

All these papers model agents as having *utilities* for alternatives, and make the assumption that the utilities of each agent sum to 1. Initiated by the work of Anshelevich et al. [2018], a large recent literature instead models agents as having *costs* for alternatives, 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, Kempe, 2020]. It turns out that the metric structure allows significantly better distortion bounds: the best distortion of deterministic rules is 3 [Gkatzelis et al., 2020] and that for probabilistic rules is between 2.0261 and $3 - 2/m$ [Charikar and Ramakrishnan, 2022, Gkatzelis et al., 2020, Kempe, 2020] (compare to $\Theta(\sqrt{m})$ distortion in the non-metric setting). Note that probabilistic rules are superior to deterministic rules in this context as well.

Fairness of single-winner voting rules has received less attention than distortion. That said, for probabilistic voting rules, fairness has been studied by in a series of papers that interpret the output distribution as a division of a budget. Historically, this was studied given approval preferences of the agents [Aziz et al., 2019, Bogomolnaia et al., 2005, Brandl et al., 2021, Duddy, 2015]. Airiau et al. [2019] studied probabilistic voting rules which take as input ranked preferences and maximize the Nash welfare (the geometric mean of agent utilities) or the egalitarian welfare (the minimum agent utility)² assuming that agent utilities are induced by their ranked preferences according to a known scoring vector. They prove that Nash-welfare-based rules satisfy SD-core, which is a weaker guarantee than the core introduced in Section 2. We note that SD-core implies no better than m -approximate of the core (for example, random dictatorship satisfies SD-core and is in m -approximate core), whereas we achieve an $O(\log m)$ -approximate core. They also proved that egalitarian-welfare-based rules satisfy an individual fairness property.

Going beyond single-winner voting, there is significant work on fairness of (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], and the proportionality degree [Skowron, 2021]. Cheng et al. [2020] proved 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] derandomized their result to prove that there always exists a committee satisfying approximate stability; we show that this derandomized result can be used to achieve $O(\sqrt{m})$

²More precisely, they studied the well-known leximin refinement of this.

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. [2016] investigated the core (which is similar to stability) for participatory budgeting, and proposed a polynomial-time algorithm for finding an outcome in the core via the so-called Lindahl equilibrium. They also pointed out connections to proportional fairness. Fain et al. [2018] studied a more general model of public goods and achieved 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 $\Delta(X)$ denote the set of all 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 $a \succeq_i a'$ to denote $\sigma_i(a) \leq \sigma_i(a')$. We refer to the collection $\vec{\sigma} = (\sigma_i)_{i \in N}$ as the *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* is a function $u : A \rightarrow \mathbb{R}_{\geq 0}$. We can extend u to also assign utility values for 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 $\vec{u} \succ \vec{\sigma}$ 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 restricted utility classes.

- $\mathcal{U}^{\text{unit-sum}}$ is the class of *unit-sum* utility functions u satisfying $\sum_{a \in A} u(a) = 1$.
- $\mathcal{U}^{\text{approval}}$ 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 $a \in A$.
- $\mathcal{U}^{\text{balanced}}$ 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 $\mathcal{U}^{\text{approval}} \subseteq \mathcal{U}^{\text{balanced}}$. Unit-sum utilities are well-studied in the distortion literature, and approval utilities occur throughout social choice theory. We introduce a new class of *balanced* utility functions, where the highest utility intensity that can be expressed is 1, which is also a lower bound on the total utility of a balanced utility function. 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 $\text{UW}(\mathbf{x}, \vec{u}) = \sum_{i \in N} u_i(\mathbf{x})$. We drop \vec{u} from the argument when it is clear from the context.

		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	$1, 1, 0$	$1, 0, 0$	$1, 1, 1$

Table 1. A preference profile with three agents and three alternatives along with two possible conforming utility profiles, where $u_1 \in \mathcal{U}^{\text{unit-sum}}$ satisfies the unit-sum condition and $u_2 \in \mathcal{U}^{\text{approval}}$ is an approval utility profile.

If one could observe the underlying utilities, an argument dating back to the work of Bentham [1789] suggests that picking the alternative maximizing the utilitarian welfare is the most efficient 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.

Definition 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 under \vec{u} and the social welfare of \mathbf{x} :

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

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} consistent with $\vec{\sigma}$.

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

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 σ with m alternatives and any number of agents.

Example 1. Take the example in Table 1 with three agents and three alternatives. For distribution $\mathbf{x} = (a_1 : 1/2, a_2 : 1/4, a_3 : 1/4)$ and utility profile \vec{u}_1 given in the table, 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, $D(\mathbf{x}, \vec{u}_1) = \frac{4/3}{13/12}$. Let us consider a different utility profile \vec{u}_2 given in the table. Now, the social welfare of \mathbf{x} is $\text{UW}(\mathbf{x}, \vec{u}_2) = 3/4 + 1/4 + 1 = 2$, but the optimal outcome given \vec{u}_2 is $\mathbf{y} = \{a_2 : 1\}$ with $\text{UW}(\mathbf{y}, \vec{u}_2) = 3$. Hence, $D(\mathbf{x}, \vec{u}_2) = 2/3$.

This is a relatively simple evaluation of the distortion of a given distribution \mathbf{x} on two possible utility profiles. Note that our eventual goal is to evaluate the distortion of \mathbf{x} under the worst-case utility profile $\vec{u} \in \mathcal{U}$, and then find the best distribution \mathbf{x}^* for which this worst case is optimized, which is significantly more complicated.

We define 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, in the absence of any cardinal preference information, no voting rule can achieve a finite distortion for \mathcal{U}^{all} because some agents (“utility monsters”) may have utilities that are orders of magnitude higher than those of other agents; thus $D_m(f, \mathcal{U}^{\text{all}}) = \infty$ for all $m \geq 2$ and all voting rules f .

Nash welfare. Distortion is usually defined with respect to the utilitarian welfare, but the same principle can be applied to other welfare functions. We will in particular study the *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 $D_m^{NW}(f, \mathcal{U})$ of a voting rule f for Nash welfare by replacing the utilitarian welfare UW in Definition 1 by NW . 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 $D_m^{NW}(f, \mathcal{U}^{\text{all}}) = D_m^{NW}(f, \mathcal{U}^{\text{unit-sum}})$.

2.2 Proportional Fairness

In addition to efficiency, we also study *fairness* of voting rules, by which intuitively we mean that every agent's preferences should be represented in the outcome. In this work, we focus on the quantitative notion of *proportional fairness* which was first proposed in communication networks [Kelly et al., 1998] but is easily adopted to social choice more generally.

Definition 2 (Proportional Fairness). Given a utility profile \vec{u} , the *proportional fairness* of a distribution over alternatives $\mathbf{x} \in \Delta(A)$ is defined as

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

This 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 $\frac{1}{n} \sum_{i \in N} \frac{u_i(\mathbf{y})}{u_i(\mathbf{x})}$ is linear in \mathbf{y} , so the maximum is in fact achieved at a distribution with singleton support.

The proportional fairness of \mathbf{x} on a preference profile $\vec{\sigma}$ for utility class \mathcal{U} is

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

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

Example 2. Take the example in Table 1. For distribution $\mathbf{x} = (a_1 : 1/2, a_2 : 1/4, a_3 : 1/4)$ and utilities \vec{u}_1 , the proportional fairness is $PF(\mathbf{x}, \vec{u}_1) = \frac{1}{3} \max\{11/3, 31/3, 13/9\} = 31/9$. For utilities \vec{u}_2 , the proportional fairness with \vec{u}_2 is $PF(\mathbf{x}, \vec{u}_2) = \frac{1}{3} \max\{7/3, 19/3, 1\} = 19/9$.

As in Example 1, this is a relatively simple evaluation of the proportional fairness of a given \mathbf{x} on two different utility profiles. Note that our eventual goal is to evaluate the proportional fairness of \mathbf{x} under the worst $\vec{u} \in \mathcal{U}$, and then find the best \mathbf{x} in this worst case setting, which is significantly more complicated.

For every utility profile \vec{u} , there exists a distribution \mathbf{x} with $PF(\mathbf{x}, \vec{u}) = 1$ [e.g., Fain et al., 2018, Sec. 2.2] which is optimal because clearly no distribution can obtain a value smaller than 1 (take $\mathbf{y} = \mathbf{x}$ in the definition of $PF(\mathbf{x}, \vec{u})$). 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 $PF(\mathbf{x}, \vec{u}) = \infty$, which we can see by taking any \mathbf{y} for which $u_i(\mathbf{y}) > 0$. Thus, a distribution whose proportional fairness value is not too high guarantees to every agent a base level of utility (in particular, it cannot completely ignore any agents' preferences).

Like the Nash welfare, proportional fairness is also scale invariant.³ Hence, we have $PF_m(f, \mathcal{U}^{\text{all}}) = PF_m(f, \mathcal{U}^{\text{unit-sum}})$ 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 $PF(\mathbf{x}, \vec{\sigma})$ and $PF_m(f)$.

³These two notions are more intimately related; see Proposition 2.

It is known that proportional fairness is a strong guarantee that implies at least two other fairness guarantees established in the literature.

Proportional fairness \Rightarrow the core. Let $\alpha \geq 1$. 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 $\frac{|S|}{|N|} \cdot u_i(\mathbf{y}) \geq \alpha \cdot u_i(\mathbf{x})$ for every agent $i \in S$ and at least one of these inequalities is strict.⁴ 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 following is a well-known relation between proportional fairness and the core.

Proposition 1. *For every $\alpha \geq 1$ and voting rule f , if $\text{PF}_m(f) \leq \alpha$, 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 $\frac{|S|}{n} u_i(\mathbf{y}) \geq \alpha \cdot u_i(\mathbf{x})$ (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 $\text{PF}_m(f) \leq \alpha$. \square

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.

Proposition 2. *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 distribution 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. \square

2.3 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, x^2 is convex and $\log x$ is concave.

Theorem 1 (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*

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

⁴Equivalently, this condition requires that there is no set $S \subseteq N$ and *partial* distribution $\mathbf{y} : A \rightarrow [0, 1]$ with $\sum_a y_a = |S|/n$ 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].

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 considering utilitarian social welfare, and voting rules with low distortion. Boutilier et al. [2015] considered unit-sum utilities, and showed that any rule must incur a distortion of at least $\Omega(\sqrt{m})$. They also constructed an intricate and artificial voting rule that achieves a distortion of $O(\sqrt{m} \log^* m)$, thus leaving a tiny gap. They also presented a more natural voting rule that achieves a distortion of $O(\sqrt{m} \log m)$, which we call the *harmonic rule* f_{HR} . It is based on the harmonic scoring rule, according to which a agent i gives $1/k$ points to the alternative that i ranks in k th position. Thus, 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 $\text{hsc}(a)$. Since the harmonic scores of all alternatives sum to nH_m where $H_m = 1 + \frac{1}{2} + \dots + \frac{1}{m}$, we can define this rule formally as

$$f_{\text{HR}}(\vec{\sigma})_a = \frac{1}{2m} + \frac{\text{hsc}(a)}{2nH_m} \quad \text{for all } a \in A.$$

In this section, we will introduce a new rule that achieves distortion $O(\sqrt{m})$ which is the optimal distortion 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 is 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 some 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 are 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. 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 simple proof of existence, following Jiang et al. [2020, Lem. 4].

Theorem 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:

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

If the bound holds, then taking an \mathbf{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

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

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 $\Pr_{a^* \sim \mathbf{y}, X \sim \mathbf{X}} [a^* \succ_i X]$ 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 $\Pr_{a^* \sim \mathbf{y}, X \sim \mathbf{X}} [a^* \succ_i X] \leq 1/(k + 1) < 1/k$.

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

Definition 4 (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 $\mathbf{x}(a) = \frac{1}{2\sqrt{m}} \cdot \Pr_{X \sim X}[a \in X] + \frac{1}{2m}$.

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 and for approval utilities. In contrast, we show in Appendix A that the harmonic rule f_{HR} achieves worse distortion for both unit-sum and (especially) approval utilities: f_{HR} achieves $\Omega(\sqrt{m \ln m})$ distortion for $\mathcal{U}^{\text{unit-sum}}$, and $\Theta(m^{2/3} \log^{1/3} m)$ distortion for $\mathcal{U}^{\text{approval}}$, demonstrating a clear separation from the performance of f_{SLR} .

Theorem 3. *On the utility class $\mathcal{U}^{\text{balanced}}$, the Stable Lottery Rule achieves $O(\sqrt{m})$ distortion:*

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

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 will use the following notation: for an alternative $a \in A$, let $\text{UW}(a) = \sum_{i \in N} u_i(a)$; for a distribution \mathbf{x} let $\text{UW}(\mathbf{x}) = \sum_{i \in N} u_i(\mathbf{x})$.

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 $\text{UW}(X) = \sum_{i \in N} u_i(X)$. Then,

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

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, $\mathbf{x}_1(a) = \frac{1}{\sqrt{m}} \cdot \Pr_{X \sim X}[a \in X]$ 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 $\text{UW}(\mathbf{x}_2) \geq \frac{n}{m}$ and so $\frac{m}{n} \cdot \text{UW}(\mathbf{x}_2) \geq 1$. Now fix an arbitrary $a^* \in A$. Then,

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

Hence,

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

Therefore, we have that

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

□

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

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

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

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 $\Omega(\sqrt{m})$; we generalize this result slightly to a broader parametric class of utility functions. Later, we provide an $\Omega(\sqrt{m})$ lower bound for the class $\mathcal{U}^{\text{approval}}$ of approval utilities. We believe it may be possible to extend our parametric result to subsume the approval result as a special case.

First, we generalize the result of Boutilier et al. [2015]. 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. [2015]. 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 further.

Theorem 4. *For any $\varepsilon \in [1/m, 1]$, the distortion of every voting rule f on the utility class \mathcal{U}^ε satisfies $D(f, \mathcal{U}) \in \Omega(\sqrt{m/\varepsilon})$.*

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

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

$a_1 > a_2 > a_3$	>	$a_4 > a_5 > a_6$	>	$a_7 > a_8 > a_9$
$a_2 > a_3 > a_1$	>	$a_5 > a_6 > a_4$	>	$a_8 > a_9 > a_7$
$a_3 > a_1 > a_2$	>	$a_6 > a_4 > a_5$	>	$a_9 > a_7 > a_8$
$a_4 > a_5 > a_6$	>	$a_7 > a_8 > a_9$	>	$a_1 > a_2 > a_3$
$a_5 > a_6 > a_4$	>	$a_8 > a_9 > a_7$	>	$a_2 > a_3 > a_1$
$a_6 > a_4 > a_5$	>	$a_9 > a_7 > a_8$	>	$a_3 > a_1 > a_2$
$a_7 > a_8 > a_9$	>	$a_1 > a_2 > a_3$	>	$a_4 > a_5 > a_6$
$a_8 > a_9 > a_7$	>	$a_2 > a_3 > a_1$	>	$a_5 > a_6 > a_4$
$a_9 > a_7 > a_8$	>	$a_3 > a_1 > a_2$	>	$a_6 > a_4 > a_5$

Fig. 1. Lower bound construction for $\mathcal{U}^{\text{approval}}$ when $m = n = 9$. Each row is a preference ranking. The red boxes are $C(A_1)$, the blue boxes are $C(A_2)$, and the green boxes are $C(A_3)$.

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

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

as needed. \square

Next, we present a lower bound for the class of approval utility functions.

Theorem 5. *For any voting rule f , we have $D(f, \mathcal{U}^{\text{approval}}) \in \Omega(\sqrt{m})$.*

PROOF. Assume \sqrt{m} is a positive integer. Divide the alternatives into equally-sized subsets $A_1, \dots, A_{\sqrt{m}}$ such that $|A_i| = \sqrt{m}$ for all $i \in \{1, \dots, \sqrt{m}\}$. For any A_i with alternatives $\{a_1, \dots, a_{\sqrt{m}}\}$, let $C(A_i)$ be a cyclic profile over A_i consisting of \sqrt{m} rankings over elements in A_i :

$$a_1 > a_2 > \dots > a_{\sqrt{m}}, \quad a_2 > a_3 > \dots > a_1, \quad \dots, \quad a_{\sqrt{m}} > a_1 > \dots > a_{\sqrt{m}-1},$$

and let $C(A_i)_j$ represent the j^{th} entry in this profile.

Now, define $C(A_1) > C(A_2) > \dots > C(A_{\sqrt{m}})$ to be a block of \sqrt{m} agents who have preferences $C(A_1)_1 > \dots > C(A_{\sqrt{m}})_1, \dots, C(A_1)_{\sqrt{m}} > \dots > C(A_{\sqrt{m}})_{\sqrt{m}}$. Using this notation, construct a profile as follows; an example profile for $n = m = 9$ is presented in Figure 1.

$$\begin{aligned} & C(A_1) > C(A_2) > \dots > C(A_{\sqrt{m}}) \\ & C(A_2) > C(A_3) > \dots > C(A_1) \\ & \vdots \\ & C(A_{\sqrt{m}}) > C(A_1) > \dots > C(A_{\sqrt{m}-1}) \end{aligned}$$

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 subset A_i such that the probability that \mathbf{x} chooses one of the alternatives in A_i is at most $1/\sqrt{m}$. Without loss of generality, assume that this is A_1 . Then, let the first \sqrt{m} agents approve the top \sqrt{m} alternatives; i.e., all alternatives in A_1 . All other agents only approve their top alternative.

Given this construction, $\text{UW}(a) = \sqrt{m}$ for all alternatives $a \in A_1$, and $\text{UW}(a') = 1$ for all alternatives $a' \in A \setminus A_1$. Therefore, for an alternative $a^* \in A_1$,

$$D(f, \mathcal{U}^{\text{approval}}) \geq \frac{UW(a^*)}{\mathbb{E}_{a \sim f} UW(a)} \geq \frac{\sqrt{m}}{(1 - \frac{1}{\sqrt{m}}) \cdot 1 + \frac{1}{\sqrt{m}} \cdot \sqrt{m}} \geq \frac{\sqrt{m}}{2} \in \Omega(\sqrt{m}).$$

□

3.4 Computation

Recall that our worst-case-optimal stable lottery rule (f_{SLR}) can be computed in polynomial time due to polynomial-time computation of a stable lottery [Cheng et al., 2020]. Hence, we turn our attention to investigating the *instance-optimal* rule. That is, we discuss how to compute, given a preference profile $\vec{\sigma}$, the instance-optimal distribution \mathbf{x} minimizing $D^{\text{UW}}(\mathbf{x}, \vec{\sigma}, \mathcal{U})$. For $\mathcal{U}^{\text{unit-sum}}$, Boutilier et al. [2015] constructed a linear program that computes such an \mathbf{x} . It is easy to see that their approach works for any utility class defined by linear inequalities, and hence it also works for $\mathcal{U}^{\text{balanced}}$.

The same method does not work for $\mathcal{U}^{\text{approval}}$ because this class is discrete. We use a different approach based on a separation oracle to find an instance-optimal distribution. Let $\vec{\sigma}$ be a preference profile. Consider a fixed distribution \mathbf{x} . First we discuss how to calculate the distortion $D^{\text{UW}}(\mathbf{x}, \vec{\sigma}, \mathcal{U}^{\text{approval}})$ of that distribution. Our task is to construct a worst-case approval utility profile \vec{u} . To do that, let us first guess which alternative $a^* \in A$ will maximize social welfare (guessing can be implemented by iterating over all options). Now, each agent $i \in N$ can be of two types: either i approves a^* and so $u_i(a^*) = 1$, or i does not and so $u_i(a^*) = 0$.

- If i approves a^* then necessarily i also approves all alternatives a with $a \succ_i a^*$. But in order to be worst-case we need to minimize $u_i(\mathbf{x})$ and thus we would set $u_i(a) = 0$ for all a with $a^* \succ_i a$. Write $r_i = \sum_{a: a \succ_i a^*} x_a$ for the value of $u_i(\mathbf{x})$ that would result under this utility function.
- If i disapproves a^* , then to minimize $u_i(\mathbf{x})$ we would set $u_i(a) = 0$ for all alternatives except for the alternative b that i ranks in top position (recalling that the definition of $\mathcal{U}^{\text{approval}}$ requires agents to approve at least one alternative). Write $s_i = x_b$ for the value of $u_i(\mathbf{x})$ that would result under this utility function.

Thus, we have narrowed down the options for the worst-case utility profile \vec{u} to two possibilities per agent. To decide which agent is of which type, first guess the number t of agents who approve a^* . Given this guess, we need to choose t agents so as to minimize $UW(\mathbf{x})$. The optimal agents to choose are the t agents with the smallest value of $r_i - s_i$. This gives a polynomial-time algorithm for computing the distortion of a given distribution \mathbf{x} .

Now consider the following linear program with exponentially many constraints:

$$\begin{aligned} & \text{maximize } \beta \\ & \text{subject to } \sum_{i \in N} u_i(\mathbf{x}) \geq \beta \cdot \max_{a \in A} \sum_{i \in N} u_i(a) \quad \text{for all } \vec{u} \in (\mathcal{U}^{\text{approval}})^n \\ & \quad \beta \geq 0, \mathbf{x} \geq 0, \sum_{a \in A} x_a = 1 \end{aligned}$$

An optimal solution \mathbf{x} to this program has instance-optimal distortion $1/\beta$. (Note that the maximum on the right-hand side of the constraints is a constant.) This program has exponentially many constraints, but note that above we constructed a polynomial-time separation oracle for it: given a candidate solution \mathbf{x} and target distortion β , we showed how to find a worst-case utility profile \vec{u} . If some constraint is violated, then the constraint corresponding to \vec{u} is violated. Hence we can use the ellipsoid method to compute an optimum \mathbf{x} . Note that this technique works for any utility class \mathcal{U} for which we can find the worst-case utility profile for a given \mathbf{x} in polynomial time.

4 PROPORTIONAL FAIRNESS

In this section, we turn our attention to proportional fairness (see Definition 2). As we noted in Section 2.2, the proportional fairness objective is scale invariant, and thus $\text{PF}_m(f, \mathcal{U}^{\text{all}}) = \text{PF}_m(f, \mathcal{U}^{\text{unit-sum}})$ 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.

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 B, 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]. We show that this rule in fact achieves the same $O(\sqrt{m})$ bound for proportional fairness. 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 upper bound the optimal proportional fairness by $O(\log m)$, which we later show to be tight. In Section 4.1, 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} . In particular, the lemma shows that the worst case for proportional fairness is always achieved by approval utilities, and so without loss of generality we can focus on them. It also follows that $\text{PF}_m(f, \mathcal{U}^{\text{all}}) = \text{PF}_m(f, \mathcal{U}^{\text{approval}})$ for all voting rules f .

Let us write $h_i(a) = \{a' : a' \succsim_i a\}$ 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 1. *Given a preference profile $\vec{\sigma}$, the proportional fairness of distribution \mathbf{x} is equal to*

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

PROOF. First, note that we can argue about the utilities of different agents separately as their contribution to $\frac{1}{n} \sum_{i \in N} \frac{u_i(\mathbf{y})}{u_i(\mathbf{x})}$, for fixed \mathbf{x} and \mathbf{y} , is independent of others.

Reduction to approval utilities. For fixed distributions \mathbf{x} and \mathbf{y} and agent i , we will show that $\frac{u_i(\mathbf{y})}{u_i(\mathbf{x})}$ is maximized for an approval utility function u_i .

Let $u_i : A \rightarrow \mathbb{R}_{\geq 0}$ be an arbitrary utility function. 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 , i.e., 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,

$$\frac{u_i(\mathbf{y})}{u_i(\mathbf{x})} = \frac{\sum_{j \in [m]} \alpha_j \vec{v}_j(\mathbf{y})}{\sum_{j \in [m]} \alpha_j \vec{v}_j(\mathbf{x})} \leq \max_{j \in [m]} \frac{\alpha_j \vec{v}_j(\mathbf{y})}{\alpha_j \vec{v}_j(\mathbf{x})} = \max_{j \in [m]} \frac{\mathbf{y}(h_i(a_j))}{\mathbf{x}(h_i(a_j))},$$

where the inequality comes from the fact that if $\frac{a}{b}, \frac{c}{d} \in \{m, M\}$ then $m \leq \frac{a+c}{b+d} \leq M$.

Simplified formulation. By Definition 2, we have $\text{PF}(\mathbf{x}, \vec{\sigma}) = \sup_{\vec{u}} \max_{a \in A} \frac{1}{n} \sum_{i \in N} \frac{u_i(a)}{u_i(\mathbf{x})}$. This is attained at a profile \vec{u} of approval utilities and at some alternative $a \in A$. In this profile, each agent $i \in N$ either approves or disapproves a . If i does not approve a , then the agent contributes a ratio of $\frac{u_i(a)}{u_i(\mathbf{x})} = 0$ which is not optimal. Hence every agent approves a . Given this, we now need to select approvals that minimize the quantity $u_i(\mathbf{x})$ for each agent. Thus, it would be suboptimal for any i to

approve any alternative ranked lower than a , as this can only increase $u_i(\mathbf{x})$. Therefore, an approval utility function maximizing the ratio $\frac{u_i(a)}{u_i(\mathbf{x})}$ for agent i approves exactly the alternatives $h_i(a)$. Thus i contributes a ratio of $\frac{1}{\mathbf{x}(h_i(a))}$ to the proportional fairness objective. Now (3) follows. \square

One helpful observation about Equation (3) is that it is convex in \mathbf{x} . This follows because $1/x$ is a convex function and taking the sum and maximum of a collection of convex functions yields a convex function.

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

Theorem 6. *There exists a voting rule which achieves proportional fairness at most $2(1 + \ln(2m))$.*

PROOF. We consider the instance-optimal voting rule which finds the distribution \mathbf{x} with optimal proportional fairness. We interpret this as the outcome of a zero-sum game, and then bounding that value in a dual game obtained by applying the minimax theorem.

Formulation as a zero-sum game. Lemma 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

$$R(\mathbf{x}, a) = \frac{1}{n} \sum_{i \in N} \frac{1}{\mathbf{x}(h_i(a))} \quad \text{and thus} \quad \text{PF}(\vec{\sigma}) = \min_{\mathbf{x} \in \Delta(A)} \max_{a \in A} R(\mathbf{x}, a). \quad (4)$$

Suppose we allow the adversary to choose *mixed* strategies, i.e., distributions over alternatives $S_A \in \Delta(A)$. Define the expected payoff of the pair (\mathbf{x}, S_A) of strategies to be $\mathbb{E}_{a \sim S_A} [R(\mathbf{x}, a)]$. This way, the game is

$$\min_{\mathbf{x} \in \Delta(A)} \max_{S_A \in \Delta(A)} \mathbb{E}_{a \sim S_A} [R(\mathbf{x}, a)]. \quad (5)$$

Note that in this formulation, the best response of the adversary can be assumed to be a pure strategy (selecting a single alternative $a \in A$) by linearity. Therefore, (4) is in fact equal to $\text{PF}(\vec{\sigma})$. Now note that $\mathbb{E}_{a \sim S_A} [R(\mathbf{x}, a)]$ is linear (and hence concave) in S_A and convex in \mathbf{x} . Therefore, by the minimax theorem (Theorem 1), we have

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

We call the game on the left-hand side 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 acts suboptimally and responds with the strategy $\bar{\mathbf{x}}$ with $\bar{\mathbf{x}}(a) = \frac{1}{2}S_A(a) + \frac{1}{2m}$ for all $a \in A$. 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. Therefore, for any preference profile $\vec{\sigma}$,

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

The last term is maximized at a distribution S_A and agent i with preference ranking of σ_i . For simplicity, 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 $T_0 = 0$. Then, the above is equal to

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

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

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

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

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

Theorem 7. *No voting rule can achieve proportional fairness smaller than $\ln(m+1)$.*

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.

Take the cyclic preference profile, where n/m agents have the preference ranking $a_j > a_{j+1} > \dots > a_m > a_1 > \dots > a_{j-1}$ for all $j \in [m]$. For simplicity, we assume $n = m$, but the proof works when n is a multiple of m .

Let \mathbf{x} be a distribution. For $i \in [m]$ and $\ell \in [m]$, define $T_i(\ell) = \sum_{k=0}^{\ell-1} \mathbf{x}(a_{i+k})$ with subscripts modulo m . Using this notation, we have $\mathbf{x}(h_i(a)) = T_i(\sigma_i(a))$ for any $a \in A$. Then,

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

Note that for each fixed ℓ , the sum of the denominators of the inner summation is $\sum_{i \in N} T_i(\ell) = \ell$ as each alternative appears in the top ℓ positions of exactly ℓ of the rankings. By the convexity of $1/x$, $\sum_{i \in N} \frac{1}{T_i(\ell)}$ is lower bounded by $m \cdot m/\ell$, which is the sum when all the $T_i(\ell)$'s are equal to ℓ/m , which is their average. Therefore, the above is lower bounded by $\sum_{\ell=1}^m 1/\ell \geq \ln(m+1)$. \square

Lower bound for Nash welfare. The same profile analyzed in Theorem 7 can be applied to distortion with Nash welfare. The difference is that instead of arguing about $\sum \frac{1}{T_i(\ell)}$, we want to find a lower bound for $\sum \log \frac{1}{T_i(\ell)}$. As $-\log(x)$ is also a convex function, the proof proceeds along the same lines and we get a lower bound of $\sum_{\ell=1}^m \log \frac{m}{\ell}$, which is equivalent to $D_m^{\text{NW}} \geq (\prod_{\ell=1}^m \frac{m}{\ell})^{1/m} = \frac{m}{m^{1/m}}$. Using Stirling's approximation, that ratio is approximately $e \cdot (\frac{1}{2\pi m})^{1/2m}$ which tends to e as $m \rightarrow \infty$.

Proposition 3. *Every voting rule has distortion at least $(\frac{m^m}{m!})^{1/m}$ with respect to the Nash welfare.*

Note that for $m = 2$, the lower bound in Proposition 3 is $\sqrt{2}$. In Appendix C, we show that this is tight by proposing a voting rule that achieves the $\sqrt{2}$ bound.

4.1 Computation

Now, we turn to the problem of computing an optimal (with respect to proportional fairness) distribution over alternatives for a given preference profile $\vec{\sigma}$. The key observations allowing polytime computation are first, convexity of $\text{PF}(\mathbf{x}, \vec{\sigma})$ in \mathbf{x} which enables the use of convex optimization methods, and second, that for a fixed distribution \mathbf{x} , $\text{PF}(\mathbf{x}, \vec{\sigma})$ can be computed in time linear to input size. To show this task can be done in polynomial time, we can use the projected subgradient descent algorithm.

Theorem 8 (Nesterov [2003], Chapter 3.2, Vishnoi [2021], Theorem 7.1). *There exists a subgradient descent-based algorithm that given a first-order oracle access to a convex function f over a bounded closed convex set Q , a number G that for all $\mathbf{x} \in Q$ and subgradients $\mathbf{g} \in \partial f(\mathbf{x})$, $\|\mathbf{g}\|_2 \leq G$, an initial point $\mathbf{x}^0 \in Q$, a number D that $\|\mathbf{x}^0 - \mathbf{x}^*\|_2 \leq D$ where $\mathbf{x}^* = \arg \min_{\mathbf{x}} f(\mathbf{x})$, and an $\varepsilon > 0$, outputs a sequence $\mathbf{x}^0, \mathbf{x}^1, \dots, \mathbf{x}^{T-1}$ such that*

$$\frac{1}{T} \sum_{i=0}^{T-1} f(\mathbf{x}^i) - f(\mathbf{x}^*) \leq \varepsilon \text{ where } T = \left(\frac{DG}{\varepsilon} \right)^2.$$

Theorem 9. *Given a preference profile $\vec{\sigma}$, a distribution over alternatives that achieves a proportional fairness of at most $\text{PF}(\vec{\sigma}) + \varepsilon$ can be computed in $\text{poly}(n, m, 1/\varepsilon)$ time.*

PROOF. We use the projected subgradient method described in Theorem 8, on $\text{PF}(\mathbf{x}, \vec{\sigma})$ which we have shown is convex in \mathbf{x} . Note that for a given $\mathbf{x} \in \Delta(A)$, the value of this function, which is $\max_{a \in A} R(\mathbf{x}, a)$, can be computed via an oracle in $\text{poly}(n, m)$ time. We choose a subset of the probability simplex such that we can guarantee bounds on the norm of the subgradients.

Let Q be the subspace of $\Delta(A)$ that for all alternatives $\{\mathbf{x}(a) \geq \alpha_a/\beta, \forall a \in A\}$, where α_a is the portion of agents who rank a as their top choice, and $\beta = 2(1 + \ln(2m))$ is the upper bound proven in Theorem 6. We will show that the optimal lottery lies in Q , and that the norm of subgradients in Q are $\text{poly}(n, m)$ size bounded. The starting point can be $\mathbf{x}_a^0 = \alpha_a, a \in A$, and $D = 2$ suffice as any two distributions in the probability simplex are in ℓ_2 distance of at most 2.

Optimality. Let \mathbf{x}^* be the optimal lottery in $\Delta(A)$. Suppose, by contradiction, that $\mathbf{x}^* \notin Q$ and there exists an alternative a with $\alpha_a > 0$ such that $\mathbf{x}(a) < \frac{\alpha_a}{\beta}$. Then,

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

However, this contradicts $\text{PF}(\vec{\sigma}) \leq \beta$. Therefore, $\mathbf{x}^* \in Q$.

Bounding Subgradients. As $R(\mathbf{x}, a)$, for a fixed $a \in A$, is differentiable and convex in \mathbf{x} , the subgradient set for the maximum of a set of convex functions, i.e. for $\max_{a \in A} R(\mathbf{x}, a)$, is the convex hull of the gradients of $R(\mathbf{x}, a)$'s that achieve the maximum at \mathbf{x} . Note that $\frac{\partial}{\partial x_{a'}} R(\mathbf{x}, a) = \sum_{i, a' \succ_i a} \frac{-1}{x(h_i(a))^2}$, and for all $i \in N$ and $a \in A$, $x(h_i(a)) \geq \min_{a' \in A} \frac{\alpha_{a'}}{\beta} \geq \frac{1}{2n(1 + \log(2m))}$. Therefore, the ℓ_∞ norm in the subgradient set is bounded by $O(n^3(\log m + 1)^2)$. This completes the criteria required by Theorem 8 to have a $\text{poly}(n, m, 1/\varepsilon)$ running time. \square

5 DISCUSSION

We prove 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 also initiate 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 $\Theta(\log m)$.

While the $O(\log m)$ upper bound on proportional fairness carries over to distortion respect to the Nash welfare, 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 $\Theta(m)$,⁵ 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 the future 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., 2021]? Can we reconcile these quantitative approaches with the classical axiomatic approach to find rules that achieve the best of both worlds?

REFERENCES

- 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.
- 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.
- 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.
- Elliot Anshelevich and John Postl. 2017. Randomized social choice functions under metric preferences. *Journal of Artificial Intelligence Research* 58 (2017), 797–827.
- 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*.
- Haris Aziz, Anna Bogomolnaia, and Hervé Moulin. 2019. Fair mixing: the case of dichotomous preferences. In *Proceedings of the 20th*. 753–781.
- Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh. 2017a. Justified representation in approval-based committee voting. *Social Choice and Welfare* 48, 2 (2017), 461–485.

⁵The 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.

- Haris Aziz, Edith Elkind, Piotr Faliszewski, Martin Lackner, and Piotr Skowron. 2017b. The Condorcet principle for multiwinner elections: from shortlisting to proportionality. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*. 84–90.
- 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)*.
- Gerdus Benade, Swaprava Nath, Ariel D. Procaccia, and Nisarg Shah. 2021. Preference elicitation for participatory budgeting. *Management Science* 67, 5 (2021), 2813–2827.
- Jeremy Bentham. 1789. *An introduction to the principles of morals and legislation*. T. Payne, London.
- Anna Bogomolnaia, Hervé Moulin, and Richard Stong. 2005. Collective choice under dichotomous preferences. *Journal of Economic Theory* 122, 2 (2005), 165–184.
- Allan Borodin, Daniel Halpern, Mohamad Latifian, and Nisarg Shah. 2022. Distortion in Voting with Top- t Preferences. <https://www.cs.toronto.edu/~nisarg/papers/top-t-distortion.pdf>.
- 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.
- 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*. 158–179.
- 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.
- Ioannis Caragiannis and Ariel D. Procaccia. 2011. Voting almost maximizes social welfare despite limited communication. *Artificial Intelligence* 175, 9-10 (2011), 1655–1671.
- Ioannis Caragiannis, Nisarg Shah, and Alexandros A. Voudouris. 2022. The Metric Distortion of Multiwinner Voting. (2022). arXiv:2201.13332.
- 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.
- 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.
- 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.
- 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.
- Conal Duddy. 2015. Fair sharing under dichotomous preferences. *Mathematical Social Sciences* 73 (2015), 1–5.
- 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) (Lecture Notes in Computer Science (LNCS))*. Springer-Verlag, 384–399.
- Brandon Fain, Kamesh Munagala, and Nisarg Shah. 2018. Fair Allocation of Indivisible Public Goods. In *Proceedings of the 19th*. 575–592. Full version arXiv:1805.03164.
- 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.
- Aris Filos-Ratsikas, Evi Micha, and Alexandros A. Voudouris. 2020. The distortion of distributed voting. *Artificial Intelligence* 286 (2020), 103343.
- 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)*. IEEE, 1427–1438.
- 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.
- 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.
- David Kempe. 2020. Communication, distortion, and randomness in metric voting. In *Proceedings of 34th AAAI Conference on Artificial Intelligence*, Vol. 34. 2087–2094.
- Martin Lackner and Piotr Skowron. 2020. Approval-Based Committee Voting. (2020). arXiv:2007.01795.
- 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.
- 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.
- Yurii Nesterov. 2003. *Introductory lectures on convex optimization: A basic course*. Vol. 87. Springer Science & Business Media.
- 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.

- 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.
- 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.
- Piotr Skowron. 2021. Proportionality degree of multiwinner rules. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*, 820–840.
- Nisheeth K Vishnoi. 2021. *Algorithms for convex optimization*. Cambridge University Press.
- John von Neumann. 1928. Zur Theorie der Gesellschaftspiele. *Math. Ann.* 100, 1 (1928), 295–320.

APPENDIX

A THE HARMONIC RULE

In this section, we provide a detailed analysis of the harmonic rule f_{HR} proposed by Boutilier et al. [2015]. They show that its distortion for unit-sum utilities is $D(f_{\text{HR}}, \mathcal{U}^{\text{unit-sum}}) = O(\sqrt{m \log m})$. First, we show that their analysis is tight (and the $\sqrt{\log m}$ factor cannot be eliminated) by proving a matching lower bound.

For each $a \in A$, let $\text{hsc}(a) := \sum_{i \in N} 1/\sigma_i(a)$ be 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 $\mathbf{x}(a) = \frac{1}{2m} + \frac{\text{hsc}(a)}{2nH_m}$.

Theorem 10. $D(f_{\text{HR}}, \mathcal{U}^{\text{unit-sum}}) \in \Omega(\sqrt{m \log m})$.

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

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

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 approval utilities. Strikingly, while the distortion of f_{HR} for unit-sum utilities is only a sublogarithmic factor worse than the best distortion, we find that its distortion for approval utilities is $\Theta(m^{2/3} \log^{1/3} m)$, which is worse than the optimal distortion of $\Theta(\sqrt{m})$ by a polynomial factor. Note that our stable lottery rule f_{SLR} achieves $\Theta(\sqrt{m})$ distortion for both unit-sum and approval utilities.

Theorem 11. *The distortion of the harmonic rule with respect to the class of approval utilities is $D(f_{\text{HR}}, \mathcal{U}^{\text{approval}}) = \Theta(m^{2/3} \log^{1/3} m)$.*

PROOF. We begin by proving the upper bound.

Upper bound. Fix an arbitrary instance. Let τ be a threshold value to be set later. Let $a^* \in \arg \max_a UW(a)$ be an optimal alternative. Consider two cases.

Case 1: Suppose $\text{hsc}(a^*) \geq \tau$. Then, $\Pr(a^*) \geq \frac{1}{2} \cdot \frac{\tau}{n \cdot H_m}$, yielding that the welfare approximation on this instance is at most $2nH_m/\tau$.

Case 2: Suppose $\text{hsc}(a^*) \leq \tau$. Let $Y = \{i \in N : u_i(a^*) = 1\}$. Note that $UW(a^*) = |Y|$. Because $\Pr(a) \geq 1/(2m)$ for each a , we have that the expected welfare under f_{HR} is at least $\frac{\sum_i U_i}{2m}$. Hence, the welfare approximation on this instance is at most $2m|Y|/(\sum_i U_i)$.

On the one hand, we can upper bound this by $2m|Y|/n$ because $U_i \geq 1$ for each i .

On the other hand, we can use the fact that

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

where the first inequality is due to the assumption of this case, the second inequality is because every $i \in Y$ ranks a^* among the first U_i positions, and the third inequality is the AM-HM inequality. Using this bound, we note that the welfare approximation in this case is also upper bounded by $2m\tau/|Y|$.

Hence, the welfare approximation in this case is upper bounded by

$$\min\left(\frac{2m|Y|}{n}, \frac{2m\tau}{|Y|}\right) \leq 2m \cdot \sqrt{\frac{\tau}{n}}.$$

Finally, combining the two cases, we can see that the distortion is at most

$$\max\left(\frac{2nH_m}{\tau}, 2m \cdot \sqrt{\frac{\tau}{n}}\right).$$

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

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

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

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

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

Finally, we show that the harmonic rule also achieves $\Theta(\sqrt{m \log m})$ proportional fairness. However, since we have shown that the best proportional fairness is $\Theta(\log m)$, this is also worse by a polynomial factor.

Theorem 12. *The proportional fairness of the harmonic rule is $\text{PF}(f_{\text{HR}}) = \Theta(\sqrt{m \log m})$.*

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 the proportional fairness objective is given by

$$\text{PF}(\mathbf{x}, \vec{u}) = \max_{a \in A} \frac{1}{n} \cdot \sum_{i \in N} \frac{u_i(a)}{u_i(\mathbf{x})} \leq \max_{a \in A} \frac{1}{n} \cdot \sum_{i \in N} \frac{1}{\mathbf{x}(h_i(a))}, \quad (6)$$

where $h_i(a) = \{b : b \succ_i a\}$ 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 (6), 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 (6), 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

B 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, achieves $O(\sqrt{m})$ proportional fairness, and therefore, $O(\sqrt{m})$ distortion with respect to the Nash welfare. Let us first formally introduce approximate stability for committees.

Definition 5 (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 [Jiang et al., 2020, Thm. 4]. However, by derandomizing the stable lottery of Cheng et al. [2020], Jiang et al. [2020] proved the following:

Theorem 13 (Jiang et al. 2020). *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 6 (c -Stable Committee Rule, c - f_{SCR}). Let X be a c -stable committee of size $k = \sqrt{m}$. The c -Stable Lottery Rule (c - f_{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 c - f_{SCR} achieves $O(\sqrt{m})$ proportional fairness when c is constant.

Theorem 14. *For constant c , proportional fairness of c - f_{SCR} is $\text{PF}(c\text{-}f_{\text{SCR}}) = O(\sqrt{m})$.*

PROOF. For constant c , consider the c - f_{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 1, we have that

$$\text{PF}(\mathbf{x}, \vec{\sigma}) = \max_{a \in A} \frac{1}{n} \sum_{i \in N} \frac{1}{\mathbf{x}(h_i(a))} = \frac{1}{n} \sum_{i \in N} \frac{1}{\mathbf{x}(h_i(a^*))}, \quad (7)$$

where $\mathbf{x}(h_i(a))$ is the probability placed on the set of alternatives $h_i(a) = \{a' : a' \succcurlyeq_i a\}$ 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 (7), 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 2. For constant c , distortion of c - f_{SCR} with respect to the Nash welfare is $D^{\text{NW}}(c\text{-}f_{\text{SCR}}, \mathcal{U}^{\text{all}}) = O(\sqrt{m})$.

We end this section by recalling that in Section 4, we improve this bound to $O(\log m)$ by using a different technique based on the minimax theorem.

C DISTORTION WITH RESPECT TO THE NASH WELFARE WITH TWO ALTERNATIVES

In this section, we analyze the optimal distortion with respect to the Nash welfare for the case of two alternatives, a_1 and a_2 . As there are only two possible rankings over 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 .

Voting Rule $f_{2\text{Nash}}$. Consider the voting rule $f_{2\text{Nash}}$, which, given any preference profile $\vec{\sigma}$, chooses $\beta(\vec{\sigma})$ satisfying

$$\frac{\ln(\beta(\vec{\sigma}))}{\ln(1 - \beta(\vec{\sigma}))} = \frac{1 - \alpha(\vec{\sigma})}{\alpha(\vec{\sigma})}. \quad (8)$$

In other words, the rule sets $\beta(\vec{\sigma}) = g^{-1}(\alpha(\vec{\sigma}))$, where $g(x) = \frac{\log(1-x)}{\log(x)+\log(1-x)}$. Note that we can disregard the extreme cases where $\alpha(\vec{\sigma}) \in \{0, 1\}$ because in such cases, there is alternative is preferred by all agents (a_1 if $\alpha(\vec{\sigma}) = 1$ and a_2 if $\alpha(\vec{\sigma}) = 0$), which the voting rule would choose.

Next, we prove that this rule achieves the optimal distortion $\sqrt{2}$ with respect to the Nash welfare.

Theorem 15. *For $m = 2$ alternatives, voting rule $f_{2\text{Nash}}$ achieves distortion $D^{\text{NW}}(f_{2\text{Nash}}, \mathcal{U}^{\text{all}}) = \sqrt{2}$ with respect to the Nash welfare, which is the best possible among all voting rules.*

PROOF. The lower bound for all voting rules proved in Proposition 3 is indeed the desired lower bound of $\sqrt{2}$ when $m = 2$. Hence, it remains to prove the upper bound for $f_{2\text{Nash}}$. Fix any preference profile $\vec{\sigma}$. For conciseness, let us write $\alpha \triangleq \alpha(\vec{\sigma})$ and $D^{\text{NW}}(\alpha) \triangleq D^{\text{NW}}(\vec{\sigma})$ as the optimal distortion on profile $\vec{\sigma}$. Note that

$$D^{\text{NW}}(\alpha) = \min_{\mathbf{x} \in \Delta(A)} \max_{\mathbf{y} \in \Delta(A), \vec{u}: \vec{u} \succ \vec{\sigma}} \left(\prod_{i \in N} \frac{u_i(\mathbf{y})}{u_i(\mathbf{x})} \right)^{1/n}. \quad (9)$$

Reduction to Approval Utilities. As shown in the proof of Lemma 1, there exist a utility profile consistent with $\vec{\sigma}$ using approval utilities which maximizes $\prod_{i \in N} \frac{u_i(\mathbf{y})}{u_i(\mathbf{x})}$, for fixed distributions \mathbf{x} and \mathbf{y} , among all utility profiles consistent with $\vec{\sigma}$. Under this worst-case utility profile, α fraction of agents i 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 (9) as

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

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 (8), which our rule $f_{2\text{Nash}}$ chooses. Further, at this β , the distortion achieved is $D^{\text{NW}}(f_{2\text{Nash}}) = \max_{\alpha \in [0,1]} D^{\text{NW}}(\alpha) = \sqrt{2}$, which is attained at $\alpha = 1/2$ (for which the optimal β is also $\beta = 1/2$). \square

From the proof of Theorem 15, it is clear that $f_{2\text{Nash}}$ is in fact an instance-optimal voting rule, not just optimal in the worst case over preference profiles.