The Distortion of Approval Voting with Runoff

Soroush Ebadian
University of Toronto
soroush@cs.toronto.edu

Mohamad Latifian
University of Toronto
latifian@cs.toronto.edu

Nisarg Shah
University of Toronto
nisarg@cs.toronto.edu

Abstract

Recent work introduces *approval with runoff* voting, in which voters cast approval ballots, two finalists are selected, and a runoff election is conducted between them to choose the final winner by majority voting. While the more common plurality with runoff voting admits only one reasonable choice of the two finalists (the two candidates with the most plurality votes), the use of approval ballots in the first stage opens up the possibility of using many reasonable ways to choose the two finalists. What is the *optimal* way to choose the two finalists?

In this work, we answer this question using the distortion framework, in which the performance of every voting system is quantitatively measured by its worst-case social welfare approximation ratio, also known as distortion. We prove that the best distortion achievable by approval voting with (majority) runoff is $\Theta(m^2)$ with deterministic finalist selection and $\Theta(m)$ with randomized finalist selection, where m is the number of candidates. This is actually worse than what simple approval voting without any runoff achieves ($\Theta(m)$ and $\Theta(\sqrt{m})$, respectively). We pinpoint the use of majority runoff in the second stage as the culprit, propose a candidate *proportional runoff* system that declares each finalist the winner with probability equal to the fraction of voters who prefer it, and analyze the extent to which it can help curb the distortion.

1 Introduction

A widely used democratic process to aggregate voters' preferences into a single winner is *two-stage voting*. In general a two-stage voting mechanism elicits some information about voters' preferences over the candidates and admit a subset of the candidates to the second stage where another round of information elicitation through voting ballots is happening. *Plurality with runoff* is the most famous two-stage voting mechanism. It is the most common method used all around the world to elect presidents.[1] In this method in the first round, each voter is asked for their top preferred candidate. If a candidate receives a simple majority (more than half) of the votes he wins the election. Otherwise, two candidates with the highest number of votes admit to the second round where a majority voting determines the final winner.

Having a runoff round is beneficial for several reasons. For instance, even if a voter's top choice does not have a chance to be the winner, his vote still counts and might effect the outcome of the election. In addition, there is a room for debates between the two rounds, and voters can deliberate on their decision. Another thing that makes voting rules with runoff interesting is that the final winner will have the support of the majority of the voters. However, plurality with runoff has several downsides that has been exhaustively studied in the literature [13]. In this regard, Delemazure et al. [13] raised the following question:

> *"... is it possible to keep the nice benefit of the two-round protocol without having to bear all the drawbacks of* plurality *at the first round?"*

[1] Plurality with runoff is used in 83 countries, while only 22 countries use the single-stage plurality. See https://en.wikipedia.org/wiki/Two-round_system.

With uni-nominal ballots, we are limited to choosing the two candidates with the highest plurality scores as the finalists. Therefore, if we want to give a positive answer to the above question we have to use a different ballot design. To this end, Delemazure et al. [13] suggest approval ballots for the first round.

Approval with runoff voting rules are not only theoretically interesting, but also they have been used in practice. In 2021, the mayoral election of St. Louis (Missouri, US) used this rule. In the first round, 44,571 voters cast approval votes for four candidates. The two candidates with the highest number of approvals admitted to the second round, and the (uni-nominal) vote of 58,237 voters determined the final winner of the election.

Using approval ballots in the first round enables the voting mechanism to apply a wider range of rules to admit two candidates to the runoff round [13]. This raises the following question: what is the *optimal* way of selecting the two finalists? Delemazure et al. [13] study this question by comparing various rules designed for approval-based committee selection (in short, ABC rules), using an axiomatic approach. However, they do not specifically discuss the quality of the final winner of the election.

Our goal is to compare different approval based voting rules with a quantitative approach. To achieve this goal we analyze the class of approval voting rules with runoff in the *implicit utilitarian framework* introduced by Procaccia and Rosenschein [25]. This framework postulates that voters have cardinal utilities for the candidates and the vote that each voter submits,stems from his utilities. That means if the voting ballot asks a voter for his ranking over the candidates the assumption of the utilitarian framework is that he ranks the candidates in the descending order of his utility for them. Similarly, with the approval ballots we assume that each voter only approves the candidates for which he has utility higher than a threshold. Our assumption is that all the voters use the same threshold, but for voters to have different thresholds is also a valid scenario. However, the latter, more general setting, does not admit any desirable distortion. Another interpretation of our setting is that the election designer decides on the threshold and asks the voters to submit their approval ballots considering this threshold.

In this framework the *social welfare* of a candidate is defined as the sum of the utilities of all the voters for her, and the optimal candidate is considered to be the one with the highest social welfare. The goal of a voting rule is to minimize the worst-case ratio between the maximum possible social welfare to the social welfare of it's outcome. This quantitative measure is called *distortion*. In approval voting with runoff, we seek *pair-selection* rules that, when followed by a majority runoff, achieve low distortion.

To capture the effect of the majority runoff on the quality of the selected winner, we investigate two other settings of (1) approval voting *without* runoff and (2) approval voting with *proportional runoff*. The former is the standard single-stage approval voting in which the winner is selected only based on the approval votes. In the latter, instead of deterministically selecting the voter with the higher number of votes, the rule selects each voter with a probability proportional to the number of the voters that select him in the runoff round. use a randomized rule two select the winner between the two finalists; each being selected with probability proportional to their votes. We find optimal distortion values in almost all of the three settings for both deterministic and randomized rules.

1.1 Our Results

We consider a voting setting with n voters and m candidates. Each voter has a utility for each candidate and following the literature [4] utilities of a voter for different candidates sums up to 1. We assume that each voter casts his approval ballot considering threshold $\tau \in [0, 1]$. This value is either selected by the election designer or is obvious from the context of the election.

We begin by analyzing the single-stage approval voting (without runoff), which we use as a benchmark to compare with the distortion values in the other with-runoff settings. We achieve (asymptotically) tight bounds on distortion for all values of $\tau \in [0, 1]$. For deterministic rules and for any $\tau \in [0, 1]$, we can achieve optimal distortion by selecting the candidate with the highest number of approvals. We show that, among all possible thresholds, $\tau = \frac{1}{m}$ gives us the optimal distortion which is $\Theta(m)$. For randomized rules, we can achieve better distortion bounds. The optimal distortion value is $\Theta(\sqrt{m})$ which is achievable at $\tau = \frac{1}{\sqrt{m}}$.

Table 1: Summary of results for the optimal distortion values of approval voting with different runoff scenarios. All upper bound results use the threshold $\tau = \frac{1}{m}$, except for the case of randomized rules with no runoffs which uses $\tau = \frac{1}{\sqrt{m}}$.

	No Runoffs	Majority	Proportional
Deterministic	$\Theta(m)$	$\Theta(m^2)$	$\Theta(m)$
Randomized	$\Theta(\sqrt{m})$	$\Theta(m)$	$O(m)$, $\Omega(m^{0.6})$

Next, we turn to approval voting with majority runoff. In this case for deterministic rules, we show that by admitting the two candidates with the highest number of approvals to the runoff, we can achieve optimal distortion for all values of $\tau \in [0, 1]$. However, in comparison to the without runoff version, distortion increases by a factor of m for all values of τ. In this setting, among all thresholds, $\tau = \frac{1}{m}$ gives us the optimal distortion of $\Theta(m^2)$. Moreover, for randomized rules, the optimal distortion increases to $\Theta(m)$. We pinpoint the use of majority runoff as the culprit.

To curb distortion in two-stage approval voting, we propose and analyze approval voting with proportional runoff. We find that by choosing the two candidates with highest number of approvals (the rule that is optimal for the voting with majority runoff) we can get back to $\Theta(m)$ distortion. This improvement also holds for all values of τ. However, for randomized rules, we show that we cannot decrease the distortion to $O(\sqrt{m})$ as any rule incurs $\Omega(m^{0.6})$ distortion.

Finally, we conduct simulations based on synthetic data, and real-world datasets of ranked preferences [23]. We evaluate the empirical performance of several ABC rules by measuring their average-case approximation-ratio of the maximum social welfare.

2 Related Work

This work extends the notion of distortion in the utilitarian framework [25], which has been extensively studied over the past decade. Caragiannis and Procaccia [12] show that the distortion of the plurality rule is $O(m^2)$. This bound was later proven to be the best possible among the deterministic ordinal voting rules [11]. Boutilier et al. [9] prove a lower bound of $\Omega(\sqrt{m})$ on the distortion of any randomized voting rule. They also present a voting rule with distortion of $O(\sqrt{m \log^* m})$. Recently, Ebadian et al. [14] closed this gap by proposing the *stable lottery rule* which achieves $\Theta(\sqrt{m})$ distortion. Borodin et al. [8] extend this result for the case that we only have access to the top-t preferred candidates of each voter.

There is also a large body of literature on the metric distortion, where voters and candidates are assumed to be embedded in a metric space [1, 2]. In this setting the distortion of any deterministic rule is known to be at least 3, and it has been proved recently that this distortion is achievable [19, 21].

In addition there are several works that consider a specific setting and give bounds on the distortion or design algorithms with good distortion in that setting [15, 18, 16, 17]. One of these works which considers a setting close to ours is the work of Pierczyński and Skowron [24] that talks about the distortion of approval voting in the metric setting.

Another line of research that our work is build on is approval voting. Voting with approval ballots has gained a lot of attention in the past decade [20, 3, 10, 22]. This type of voting ballots is useful in different settings, for instance Benade et al. [6] consider the distortion of threshold approval votes in participatory budgeting. Moreover, the approval voting with runoff rule which we are investigating was introduced by Sanver [26] and further investigated by Delemazure et al. [13].

3 Preliminaries

For $t \in \mathbb{N}$, define $[t] = \{1, 2, \ldots, t\}$, and for a finite set S, define $\Delta(S)$ to be the set of all probability distributions over S.

A single-winner election consists of a set $N = [n]$ of n voters, and a set $C = \{c_1, c_2, \ldots, c_m\}$ of m candidates. We assume that each voter $i \in N$ has a utility function $u_i : C \to \mathbb{R}_{\geqslant 0}$ over the candidates, where $u_i(c)$ is her utility for candidate c. For $x \in \Delta(C)$ we define $u_i(x) = \mathrm{E}_{c \sim x}[u_i(c)]$. Following the literature [5], we work with unit-sum utility functions, where $\sum_{c \in C} u_i(c) = 1$ for each $i \in N$.

Let $\vec{u} = (u_1, u_2, \ldots, u_n)$ be the utility profile. Define the social welfare of a candidate $c \in C$ to be $\mathrm{SW}(c, \vec{u}) = \sum_{i \in N} u_i(c)$. The goal of the election is to select a candidate with the maximum social welfare. The challenge is that we do not have access to the utilities, and eliciting them directly would place undue cognitive burden on the voters. Hence, the election designer designs a voting system, which operates in one or more sequential stages. In each stage, a voting ballot elicits some partial information about voter utilities, and the design of the ballot may depend on the information collected in previous stages.

We are primarily interested in two voting systems: (single-stage) approval voting and (two-stage) approval voting with runoff.

3.1 Approval Voting

Single-stage approval voting is conducted using approval ballots, where each voter approves an unranked subset of the candidates. We model this using a threshold $\tau \in [0, 1]$ such that each voter approves all the candidates for which she has utility at least τ. When τ is explicitly set by the designer, these is known as threshold approval voting [7], but this can also be used as a model for the voters converting utility functions to approval votes. We refer to such ballots as τ-approval ballots.

Let $\text{ß}_i \subseteq C$ be the approval vote of voter i, and $\vec{\pi} = (\text{ß}_1, \text{ß}_2, \ldots, \text{ß}_n)$ be the approval profile. A (randomized) approval voting rule $f : \left(2^C\right)^n \to \Delta(C)$ maps an approval profile to a distribution over the candidates. We say that f is deterministic if it always output a distribution with singleton support.

3.2 Approval Voting with Runoff

Approval voting with runoff is a two-stage election system. In the first stage, a pair of candidates are selected as *finalists* based on approval ballots submitted by the voters, and in the second stage, an election (called a runoff) is conducted between the two finalists, where each voter specifies whom she prefers more and a runoff rule is used to select the winner. To ensure that voters break ties between equal-utility finalists consistently in the runoff election, we assume that each voter i has a ranking $\text{œ}_i : [m] \to C$, where $\text{œ}_i(j)$ refers to the j-th most preferred candidate of voter i, that is consistent with her utility function, i.e., $u_i(\text{œ}_i(j)) \geqslant u_i(\text{œ}_i(j'))$ for all $j < j'$. In the runoff election between any two finalists, each voter i votes for the one who is higher up in her ranking œ_i. Let $\vec{\text{œ}} = (\text{œ}_1, \text{œ}_2, \ldots, \text{œ}_n)$ denote the profile of these rankings.

Formally, in the first stage, a (randomized) approval-based pair-selection rule $f_1 : \left(2^C\right)^n \to \Delta\left(C^2\right)$ takes as input an approval profile and returns a distribution over pairs of candidates. In the second stage, a runoff is conducted between a pair of candidates (finalists) sampled from this distribution. For notational simplicity, we say that a runoff rule f_2 takes as input the distribution over pairs of candidates from the first stage $f_1(\vec{\pi})$ and the ranking profile $\vec{\text{œ}}$ and returns a distribution over the candidates. Together, we can denote an entire system of approval voting with runoff as a voting rule $f_1 \circ f_2(\vec{\pi}, \vec{\text{œ}}) = f_2(f_1(\vec{\pi}), \vec{\text{œ}})$. In this work consider two runoff rules:

- *Majority Runoff (*maj*).* This is the canonical runoff method, in which, once a pair of finalists is sampled from the distribution returned in the first stage, the finalist preferred by a majority of the voters is selected as the winner (deterministically).[2]

[2]In case of a tie, where each finalist is preferred by exactly half of the voters, either can be selected as the winner and our proofs continue to work.

- *Proportional Runoff* (prop). This is a novel runoff method that we introduce and study, in which, once a pair of finalists is sampled from the distribution returned in the first stage, each finalist is selected with probability equal to the fraction of voters who prefer it. Note that unlike majority runoff, proportional runoff uses randomization to select the winner among the two finalists.

3.3 Distortion in Approval Voting

As stated earlier, our goal is to select a candidate with high social welfare. Since the voting systems defined above have access to only partial information about voter utilities, we use the notion of *distortion* proposed by Procaccia and Rosenschein [25] to quantify how well a voting rule f maximizes social welfare.

Define the distortion of a distribution over candidates $x \in \Delta(C)$ on a utility profile $\vec{u}$ as:

$$\mathsf{dist}(x, \vec{u}) = \frac{\max_{c \in C} \mathsf{SW}(c, \vec{u})}{\mathsf{SW}(x, \vec{u})}.$$

For a single-stage approval voting rule f, define its distortion for utility profile $\vec{u}$ with respect to threshold τ as $\mathsf{dist}_\tau(f, \vec{u}) = \mathsf{dist}(f(\vec{\pi}), \vec{u})$, where $\vec{\pi}$ is the approval profile induced by utility profile $\vec{u}$ under threshold τ. The (overall) distortion of f with respect to τ is $\mathsf{dist}_\tau(f) = \sup_{\vec{u}} \mathsf{dist}_\tau(f, \vec{u})$.

For a pair selection rule f_1 and runoff rule $f_2 \in \{\mathsf{maj}, \mathsf{prop}\}$ in two-stage approval voting with runoff, define its distortion for a consistent pair of utility profile $\vec{u}$ and ranking profile $\vec{œ}$ with respect to threshold τ as $\mathsf{dist}_\tau(f_2 \circ f_1, \vec{u}, \vec{œ}) = \mathsf{dist}(f_2 \circ f_1(\vec{\pi}, \vec{œ}), \vec{u})$, where $\vec{\pi}$ is again the approval profile induced by utility profile $\vec{u}$ under threshold τ. The (overall) distortion of f with respect to τ is $\mathsf{dist}_\tau(f_2 \circ f_1) = \sup_{(\vec{u}, \vec{œ})} \mathsf{dist}_\tau(f_2 \circ f_1, \vec{u}, \vec{œ})$, where the supremum is over consistent pairs of utility profiles and ranking profiles.

4 Approval Voting

We begin with an analysis of the simpler (single-stage) approval voting, which we use as a benchmark to compare to in our analysis of approval voting with runoff in the next section.

4.1 Deterministic Rules

First, we analyze the case where we want to select a winner *deterministically* based on τ-approval ballots. The next result shows that when τ is too large, none of the voters might approve any candidate, and a deterministic winner selected without any information about voter utilities might incur unbounded distortion.

Theorem 1. *For $\tau > \frac{1}{m-1}$, any deterministic rule with τ-approval ballots incurs unbounded distortion.*

Proof. Suppose none of the voters approve any candidate and a deterministic voting rule selects candidate c as the winner. We construct the underlying utility profile $\vec{u}$ in which every voter i has utility $u_i(c') = \frac{1}{m-1}$ for every candidate $c' \neq c$, and $u_i(c) = 0$. Note that this is consistent with empty τ-approval ballots for $\tau > \frac{1}{m-1}$. In this case, $\mathsf{SW}(c, \vec{u}) = 0$ whereas $\mathsf{SW}(c', \vec{u}) = \frac{n}{m-1} > 0$ for any other candidate c', resulting in unbounded distortion. ☐

Note that when $\tau \leqslant \frac{1}{m}$, every voter must approve at least one candidate, as $\max_c u_i(c) \geqslant \frac{1}{m} \sum_c u_i(c) = \frac{1}{m}$. Hence, the pessimistic scenario above, in which none of the voters approving any candidate, can no longer arise. The following result shows an upper bound we can place on the distortion in this case, using the most straightforward voting rule that selects the most approved candidate as the winner. We later show that this is optimal up to a constant factor.

Theorem 2. *For $\tau \leqslant \frac{1}{m}$, selecting the most approved candidate based on τ-approval ballots achieves a distortion of at most $\frac{1}{\tau} + m - 1$.*

Proof. Let $\vec{u}$ be a utility profile and $\vec{\pi}$ be the induced τ-approval profile. Let c be the most approved candidate with q approvals. Then, $\text{SW}(c, \vec{u}) \geqslant q \cdot \tau$. Since voter utilities are unit-sum and $\tau \leqslant \frac{1}{m}$, each voter approves at least one candidate. Hence, $q \geqslant \frac{n}{m}$.

Let c^* be an optimal candidate with the maximum social welfare under $\vec{u}$. Then, $\text{SW}(c^*, \vec{u}) \leqslant q \cdot 1 + (n-q) \cdot \tau$, since at most q voters who approve c^* can have utility at most 1 for it and the rest have utility less than τ for it. Thus, the distortion is bounded by

$$\frac{q + (n-q)\tau}{q \cdot \tau} \leqslant \frac{1}{\tau} + \frac{n}{q} - 1 \leqslant \frac{1}{\tau} + m - 1,$$

where, in the last inequality, we used the fact that $q \geqslant \frac{n}{m}$. □

We now prove that this upper bound is tight for deterministic rules up to a constant factor. Note that for $\tau \leqslant 1/m$ and $m \geqslant 2$, $1/\tau - 1 \geqslant (1/3)(1/\tau + m - 1)$.

Theorem 3. *For $m \geqslant 3$ and $\tau \leqslant \frac{1}{m-1}$, any deterministic rule with τ-approval ballots incurs a distortion of at least $\frac{1}{\tau} - 1$.*

Proof. Suppose that for each $i \in [m]$, there are n/m voters who approve candidates c_i and $c_{(i \bmod m)+1}$. Every voter has utility $1-\tau$ for their top choice, τ for their second choice, and 0 for the rest.

Let c be the candidate selected by a deterministic rule. Let the $2n/m$ voters approving c have utility τ for c and utility $1-\tau$ for their other approved candidate. Pick any candidate $c^* \neq c$, and let the $2n/m$ voters approving c^* have utility $1-\tau$ for c^* and τ for their other approved candidate.

Under this utility profile $\vec{u}$, $\text{SW}(c, \vec{u}) = \tau \cdot 2n/m$ whereas $\text{SW}(c^*, \vec{u}) = (1-\tau) \cdot 2n/m$. Hence, the distortion is at least

$$\frac{(1-\tau) \cdot 2n/m}{\tau \cdot 2n/m} = \frac{1}{\tau} - 1.$$ □

We conclude by noting the distortion achievable by deterministic rules in single-stage approval voting (along with the optimal choice of τ) based on the above results.

Corollary 1. *For deterministic rules with τ-approval ballots, the optimal distortion is $\Theta(m)$, which can be achieved by at $\tau = \frac{1}{m}$ by selecting the most approved candidate.*

Proof. By Theorem 1, distortion is unbounded when $\tau > \frac{1}{m-1}$, and for $\tau \leqslant \frac{1}{m-1}$, we know from Theorem 3 that distortion is lower bounded by $\frac{1}{\tau} - 1$, which is minimized at $\tau = \frac{1}{m-1}$. Hence, distortion is at least $m-2$. Furthermore, Theorem 2 shows that at $\tau = \frac{1}{m}$, selecting the most approved candidate achieves a distortion of at most $2m-1$. Therefore, the optimal distortion value is $\Theta(m)$ and can be achieved at $\tau = \frac{1}{m}$. □

4.2 Randomized Rules

Next, we turn to randomized rules for single-stage approval voting. Recall that when $\tau > \frac{1}{m-1}$, deterministic rules incur unbounded distortion due to the possibility of none of the voters approving any candidate. In contrast, we show that with randomized rules, better distortion bounds can be achieved precisely by using such high values of τ. Our starting point is the simple observation that selecting a candidate uniformly at random already achieves a distortion of m regardless of τ and the approval profile because the optimal candidate is selected with a probability of $\frac{1}{m}$. We show that mixing uniformly random selection with the deterministic strategy of selecting the most approved candidate achieves a significantly better distortion, as low as $O(\sqrt{m})$ when $\tau = 1/\sqrt{m}$.

Theorem 4. *For $\tau \in [\frac{1}{m}, 1]$, given τ-approval ballots, the randomized rule f that,*

1. *with probability 1/2, returns the most approved candidate, and*

2. *with probability* 1/2*, returns a candidate uniformly at random,*

achieves a distortion of at most $2(\frac{1}{\tau} + m\tau)$.

Proof. Consider any utility profile $\vec{u}$ and induced τ-approval profile $\vec{\pi}$. Let c be the most approved candidate with q approvals. Hence, $\text{SW}(c, \vec{u}) \geqslant q\tau$. It is easy to see that uniformly random selection achieves an expected social welfare of $\frac{n}{m}$. Hence, the expected social welfare under f is $\text{SW}(f(\vec{\pi}), \vec{u}) \geqslant \frac{1}{2}\left(q\tau + \frac{n}{m}\right)$. Let c^* be an optimal candidate with the maximum social welfare. Note that each of at most q voters who approve c^* has utility at most 1 for it, while the rest have utility at most τ for it. Hence, $\text{SW}(c^*, \vec{u}) \leqslant q \cdot 1 + (n-q)\tau$. Therefore,

$$\mathsf{dist}_\tau(f) \leqslant \frac{q + (n-q)\tau}{1/2\,(q\tau + n/m)} \leqslant 2\Big(\frac{1}{\tau} + m\tau\Big). \qquad \square$$

By setting $\tau = \frac{1}{\sqrt{m}}$, the theorem above yields a distortion of $O(\sqrt{m})$, which improves upon the best distortion attainable with deterministic rules.

It is worth mentioning that while using uniformly random selection with 1/2 probability may seem unjustifiable, changing the 1/2 to any (ever so small) constant probability retains the distortion guarantees of Theorem 4 asymptotically. Furthermore, if all voters prefer one candidate to another, shifting any probability mass placed on the latter (Pareto dominated) candidate to the former can only improve distortion.

Now, we show the tightness of the bounds given in Theorem 4.

Theorem 5. *Any randomized rule with τ-approval ballots incurs a distortion of at least $\frac{m}{2}$ for $\tau \leqslant \frac{1}{2m}$, and $\frac{1}{8}\big(\frac{1}{\tau} + m\tau\big)$ for $\tau \geqslant \frac{1}{2m}$.*

Proof. First, we prove the result for $\tau \geqslant 1/\sqrt{m}$. In this case, since $1/\tau \leqslant m\tau$, it is sufficient to prove a lower bound of $\frac{1}{4}m\tau$.

Consider an approval profile $\vec{\pi}$ in which for each $i \in [m]$, there are n/m voters approve only candidate c_i. Let f be any randomized rule. There must exist a candidate c that is selected under $f(\vec{\pi})$ with probability at most $\frac{1}{m}$. Now, consider a consistent utility profile $\vec{u}$, in which the n/m voters approving c have utility 1 for it and 0 for every other candidate; while every other voter has utility τ for her approved candidate, $\tau/2$ for c, and $1 - 3\tau/2$ divided equally between the remaining candidates. In this case, $\text{SW}(c, \vec{u}) \geqslant n\tau/2$, whereas $\text{SW}(c', \vec{u}) \leqslant n/(m-1)$ for every $c' \neq c$ because a total utility of at most n is divided equally between the remaining candidates due to symmetry. Hence,

$$\begin{aligned}\mathsf{dist}_\tau(f) &\geqslant \frac{\text{SW}(c, \vec{u})}{\frac{1}{m}\,\text{SW}(c, \vec{u}) + \left(1 - \frac{1}{m}\right)\frac{n}{m-1}} \\ &\geqslant \frac{1}{\frac{1}{m} + \frac{2}{(m-1)\tau}} \geqslant (m-1)\tau \geqslant \frac{1}{4}m\tau.\end{aligned}$$

Next, we consider $\frac{1}{\sqrt{m}} \geqslant \tau \geqslant \frac{1}{2m}$. Here, since $1/\tau \geqslant m\tau$, it is sufficient to prove a lower bound of $\frac{1}{4\tau}$. Consider an approval profile $\vec{\pi}$ in which every voter approves the set of candidates $A_\tau = \{1, 2, \dots, \lceil \frac{1}{2\tau} \rceil\}$. Let c be the candidate among A_τ that is selected with probability at most $\frac{1}{|A_\tau|} \leqslant 2\tau$ under $f(\vec{\pi})$. Consider a consistent utility profile $\vec{u}$ in which every voter has utility τ for every candidate in $A_\tau \setminus \{c\}$ (this is less than $1/2$ in total), and the remaining utility of at least $1/2$ for c. Hence, $\text{SW}(c, \vec{u}) \geqslant n/2$; for all $c' \in A_\tau \setminus \{c\}$, $\text{SW}(c', \vec{u}) = n\tau$; and, the social welfare of every other candidate is 0. Since c is selected with probability at most 2τ by f,

$$\mathsf{dist}_\tau(f) \geqslant \frac{\text{SW}(c, \vec{u})}{2\tau \cdot \text{SW}(c, \vec{u}) + (1-2\tau) \cdot n\tau} \geqslant \frac{1}{2\tau + 2\tau} \geqslant \frac{1}{4\tau}.$$

At $\tau = \frac{1}{2m}$, the lower bound is equal to $\frac{m}{2}$. In the adversarial profile, every voter approves all the candidates, and has utility $\frac{1}{2} + \frac{1}{2m}$ for her top choice and $\frac{1}{2m}$ for every other candidate. It is easy to see that this lower bound, via the same example, continues to hold even when $\tau < \frac{1}{2m}$. $\square$

We conclude by noting the optimal distortion attainable using randomized rules, which follows from Theorems 4 and 5.

Corollary 2. *For randomized rules with τ-approval ballots, the optimal distortion value is $\Theta(\sqrt{m})$, which is achievable at $\tau = \frac{1}{\sqrt{m}}$.*

5 Approval Voting with Majority Runoff

In this section, we show how the distortion worsens (i.e., increases) for both deterministic and randomized rules when a majority runoff is added to the process.

It is worth remarking that additional information is gained during the runoff stage, which cannot necessarily be known from the approval ballots in the first stage. One might wonder whether this additional information can help select a better candidate; we prove that this is not true, at least in the worst case. The reason is that the imposition of majority runoff constrains the process; e.g., a Condorcet loser — a candidate preferred to any other candidate by only a minority of voters — can never be selected as the final winner. The following example shows that this already incurs a distortion of $\Omega(m)$.

Example 1 (Valuable Condorcet Loser)**.** Consider the example in Table 2, in which $\frac{n}{2} - 1$ voters have a utility of 1 for candidate c_1, while the other $\frac{n}{2} + 1$ voters have zero utility for c_1.

Voter	$u_i(c_1)$	$u_i(c'), \forall c' \neq c_1$
$i \in \{1, \ldots, n/2 - 1\}$	1	0
$i \in \{n/2, \ldots, n\}$	0	$1/(m-1)$

Table 2: Utility profile of an instance where the Condorcet loser c_1 has high social welfare compared to others.

Under this utility profile $\vec{u}$, the social welfare of c_1 is $\mathrm{SW}(c_1, \vec{u}) = \frac{n}{2} - 1$, whereas $\mathrm{SW}(c', \vec{u}) = \frac{n/2+1}{m-1}$ for every $c' \neq c_1$. However, c_1 is the Condorcet loser because $\frac{n}{2} + 1$ voters prefer every other candidate to c_1, so it must be selected with zero probability under approval voting with majority runoff. Hence, the distortion of any (even randomized) rule is at least $\frac{n/2-1}{(n/2+1)/(m-1)} = \Omega(m)$.

We have shown the following.

Theorem 6. *In approval voting with majority runoff, any (even randomized) rule incurs a distortion of $\Omega(m)$, even if the rule is given access to exact voter utilities.*

In the rest of the section, we show that the optimal distortion for deterministic rules is actually worse ($\Theta(m^2)$), while this lower bound is tight for randomized rules. In both cases, the optimal distortion is achieved at $\tau = \frac{1}{m}$, in contrast to the case of no runoff.

5.1 Deterministic Rules

Similar to Theorem 1, we argue that using a large value of τ can lead to high distortion for deterministic rules, as voters may not approve any candidates. Due to space limits, we defer the proof to Appendix A.1.

Theorem 7. *In approval voting with majority runoff, for $\tau > \frac{1}{m-2}$, any deterministic rule incurs an unbounded distortion.*

In contrast, when $\tau \leqslant \frac{1}{m}$, each voter approves at least one candidate. A reasonable rule in this case returns the two candidates with the highest numbers of approvals, breaking ties arbitrarily. We analyze the distortion of this rule in the following lemma, and later prove that this is optimal up to a constant factor.

Theorem 8. *In approval voting with majority runoff, for $\tau \leqslant \frac{1}{m}$, selecting the two candidates with the highest numbers of approvals achieves a distortion of at most $\frac{2m}{\tau}$.*

Proof. Consider any instance with utility profile $\vec{u}$ and approval profile $\vec{\pi}$. Let c_1 and c_2 be the two most approved candidates with q_1 and q_2 approvals, respectively, with $q_1 \geqslant q_2$. If c_1 wins the runoff, then by Theorem 2 distortion is bounded by $\frac{1}{\tau} + m - 1$ which is better than the sought $\frac{2m}{\tau}$ bound.

Next, suppose c_2 wins the runoff. Note that c_2 must be preferred to c_1 by at least $\frac{n}{2}$ voters. Further, each such voter approves at least her most preferred candidate (which is not c_1) due to $\tau \leqslant \frac{1}{m}$. Hence, candidates in $C \setminus \{c_1\}$ have a total of at least $\frac{n}{2}$ approvals. Since c_2 has the highest number of approvals among such candidates, $q_2 \geqslant \frac{n/2}{m-1} \geqslant \frac{n}{2m}$. Thus, $\text{SW}(c_2, \vec{u}) \geqslant \frac{n\tau}{2m}$. Furthermore, as $\text{SW}(c, \vec{u}) \leqslant n - \text{SW}(c_2, \vec{u})$ for any candidate $c \neq c_2$, we have that the distortion of f is at most

$$\frac{n - \text{SW}(c_2, \vec{u})}{\text{SW}(c_2, \vec{u})} \leqslant \frac{n}{n\tau/(2m)} - 1 = \frac{2m}{\tau} - 1. \qquad \square$$

To show the optimality of the distortion bound proven above, we can utilize Example 1 and the example presented in Theorem 3. The idea is to have a Condorcet loser that is the top rank of $\frac{n}{2} - 1$ voters with utility of 1, which cannot win the election. The approval profile of the rest of the agents is as constructed in Theorem 3. In Theorem 3, we show the winner has a social welfare of at most $\frac{n}{m} \cdot \tau$. The social welfare of the Condorcet loser in this new example is $\frac{n}{2} - 1$. Therefore, we obtain a distortion lower bound of $\frac{n/2-1}{n/m \cdot \tau} \geqslant \frac{m}{2\tau} - 1$. We defer the complete proof of Theorem 9 to Appendix A.2.

Theorem 9. *For $\tau \leqslant \frac{1}{m-2}$, in approval voting with majority runoff, any deterministic rule incurs a distortion of at least $\frac{m}{2\tau} - 1$.*

We conclude by noting the distortion achievable by deterministic rules in approval voting with majority runoff (along with the optimal choice of τ) based on the above results (Theorems 7 to 9).

Corollary 3. *In approval voting with majority runoff, the optimal distortion value for deterministic rules is $\Theta(m^2)$ which can be achieved at $\tau = \frac{1}{m}$.*

5.2 Randomized Rules

Next, we turn to randomized rules for approval voting with majority runoff. Recall that Theorem 6 shows a distortion better than $O(m)$ is infeasible, even with randomization. We present an upper bound of $O(m)$ that shows the optimal distortion value for randomized rules is $\Theta(m)$. Moreover, similar to the case of deterministic rules, the outcome of adding a majority runoff to approval voting is an increase in distortion of randomized rules from $\Theta(\sqrt{m})$ to $\Theta(m)$.

The following is a technical lemma showing an upper bound on the social welfare of a Condorcet loser (if existent), which is useful for proving our upper bound. The proof is provided in Appendix A.3.

Lemma 1. *If a candidate c is the Condorcet loser in an instance, then $\text{SW}(c) \leqslant \frac{n}{2} + \frac{n}{2m}$.*

Now, we are ready to present the randomized rule that achieves the optimal distortion when a majority runoff is used. The idea is to always include the most approved candidate as a finalist, and for the second finalist, we show that mixing uniformly random selection with the deterministic strategy of selecting the second most approved candidate, achieves a distortion of $O(m)$ at $\tau = \frac{1}{m}$. Due to space limits, we defer the proof to Appendix A.4.

Theorem 10. *For $\tau = \frac{1}{m}$, let c_1 and c_2 be the two candidates with highest number of approvals, then the randomized rule f that*

1. *with probability $1/2$, selects (c_1, c_2) as finalists,*

2. *and with probability 1/2, selects the pair (c_1, c') with a random candidate $c' \in C \setminus \{c_1, c_2\}$ as finalists,*

achieves a distortion of at most $4m$.

We conclude by noting the optimal distortion attainable using randomized rules in approval voting with majority runoff.

Corollary 4. *In approval voting with majority runoff, the optimal distortion value of randomized rules is $\Theta(m)$ which can be achievable at $\tau = \frac{1}{\sqrt{m}}$.*

6 Approval Voting with Proportional Runoff

In the prior sections, we analyze distortion of approval voting with and without majority runoff. We observe that addition of majority runoff increases distortion for both deterministic and randomize rules. To find a middle-ground and curb distortion, we propose and analyze using another runoff method, proportional runoff.

6.1 Deterministic Rules

We show that by using proportional runoff in approval voting can decrease the distortion of deterministic rules to $\Theta(m)$ compared to the distortion of $\Theta(m^2)$ when majority runoff is used.

Theorem 7 shows that for $\tau > \frac{1}{m-2}$, there exists an example such that the pair of candidates picked by any deterministic pair selection have zero social welfare and, hence, unbounded distortion.

Corollary 5 (Theorem 7)**.** *In approval voting with proportional runoff, for $\tau > \frac{1}{m-2}$, any deterministic rule incurs unbounded distortion.*

Recall that when $\tau \leqslant \frac{1}{m}$, every voter must approve at least one candidate. The following result analyzes the pair selection that elects the two most approved alternatives as finalists (breaking tie arbitrarily). We later prove that this is optimal up to a constant factor. In Theorem 8, we show that the same method achieves the optimal distortion with majority runoff.

Theorem 11. *In approval voting with proportional runoff, for $\tau \leqslant \frac{1}{m}$, selecting the two candidates with the highest number of approvals achieves a distortion of at most $8\big(\frac{1}{\tau} + m\big)$.*

Proof. Consider any instance with utility profile $\vec{u}$ and approval profile $\vec{\pi}$. Suppose c_1 and c_2 are the two most approved candidates with q_1 and q_2 approvals respectively. If the majority of voters prefer c_1 to c_2, then c_1 is selected with probability at least 1/2, and, by Theorem 2, distortion is bounded by $2 \cdot (\frac{1}{\tau} + m - 1)$. Otherwise, c_2 wins the pairwise comparison and is selected with probability at least 1/2.

Among the q_1 voters who approve c_1, either half of them have $c_1 \succ c_2$ and $\Pr[c_1] \geqslant \frac{q_1/2}{n}$, or at least half of them have $c_2 \succ c_1$ and approve both c_2 and c_1. In the latter, $q_2 \geqslant \frac{q_1}{2}$. Therefore, with probability at least 1/2, we select a candidate with $\frac{q_1}{2}$ (half of the maximum) approvals. Following the proof of Theorem 2, distortion would be bounded by $4 \cdot (\frac{1}{\tau} + m - 1)$. The only remaining case, is when $\Pr[c_1] \geqslant \frac{q_1/2}{n}$.

- Case $c^* = c_1$. Then, from $q_1 \geqslant \frac{n}{m}$, we have $\Pr[c_1] = \Pr[c^*] \geqslant \frac{q_1/2}{n} \geqslant \frac{1}{2m}$. Hence, distortion is bounded by $2m$.
- Case $c^* \neq c_1$. Then, c^* has at most q_2 approvals, and

$$
\begin{aligned}
\mathrm{E}[\mathsf{SW}(\mathsf{prop} \circ f, \vec{u})] &\geqslant (\Pr[c_1] \cdot q_1 + \Pr[c_2] \cdot q_2) \cdot \tau \\
&\geqslant \left(\frac{q_1}{2n} \cdot q_1 + \frac{1}{2} \cdot \frac{n - q_1}{m} \right) \tau.
\end{aligned}
$$

If $q_1 \geqslant \frac{n}{\sqrt{2m}}$, we have $\frac{q_1 \cdot q_1}{n} \geqslant \frac{n}{2m}$. Otherwise, as $\frac{n}{\sqrt{2m}} \leqslant \frac{n}{2}$ (holds for $m \geqslant 2$), we have $\frac{n-q_1}{m} \geqslant \frac{n}{2m}$. Hence, $\mathrm{E}[\mathsf{SW}(\mathsf{prop} \circ f, \vec{u})] \geqslant \frac{n}{4m} \cdot \tau$. This bound combined with $\mathrm{E}[\mathsf{SW}(f)] \geqslant \Pr[c_2] \cdot q_2 \tau$ results in

$$\mathrm{E}[\mathsf{SW}(\mathsf{prop} \circ f, \vec{u})] \geqslant \frac{1}{2} \left(\Pr[c_2] \cdot q_2 \tau + \frac{n}{4m} \cdot \tau \right).$$

Furthermore, $\mathsf{SW}(c^*, \vec{u}) \leqslant q_2 + (n - q_2)\tau$ (approvals with utility 1 and non-approvals with utility at most τ). Therefore, distortion is bounded by

$$\frac{\mathsf{SW}(c^*, \vec{u})}{\mathrm{E}[\mathsf{SW}(\mathsf{prop} \circ f, \vec{u})]} \leqslant \frac{q_2 + (n - q_2)\tau}{\frac{1}{2} \left(\Pr[c_2] \cdot q_2 \tau + \frac{n}{4m} \tau \right)}$$
$$\leqslant 2 \cdot \frac{q_2 + n\tau}{\frac{1}{2} \cdot q_2 \tau + \frac{n}{4m} \tau} \leqslant 8 \left(\frac{1}{\tau} + m \right).$$

□

We now prove that this upper bound is tight for deterministic rules up to a constant factor. Note that for $\tau \leqslant \frac{1}{m}$, $\frac{1}{\tau} + m \leqslant \frac{2}{\tau}$. As the lower bound construction is similar to Theorem 3, we defer the proof to Appendix A.5.

Theorem 12. *In approval voting with proportional runoff, for $\tau \leqslant \frac{1}{m-2}$, any deterministic incurs a distortion of at least $\frac{1}{2} \left(\frac{1}{\tau} - 1 \right)$.*

We conclude by noting the optimal distortion for deterministic rules with proportional runoff.

Corollary 6. *In approval voting with proportional runoff, the optimal distortion value for deterministic rules is $\Theta(m)$ which can be achieved at $\tau = \frac{1}{m}$.*

6.2 Randomized Rules

Now, we turn to randomized rules in approval voting with proportional runoff. The improvement in distortion of deterministic rules with proportional runoff compared to majority runoff, hints that distortion of randomized rules should also improve from $\Theta(m)$ when majority runoff is used. While we were unsuccessful in finding upper bounds better than $O(m)$, we prove that distortion in this setting is $\omega(\sqrt{m})$; more specifically, we prove that it is $\Omega(m^{0.6})$, and we leave this gap in the optimal distortion value (between $\Omega(m^{0.6})$ and $O(m)$), as an open problem.

The following observation helps to apply some known results in previous settings to the setting with proportional runoff.

Proposition 1. *Let f be a pair selection rule. Then, for all $\tau \in [0, 1]$ and m, we have $\frac{1}{2} \mathsf{dist}_\tau(m) \leqslant \mathsf{dist}_\tau(\mathsf{prop} \circ f) \leqslant 2 \cdot \mathsf{dist}_\tau(\mathsf{maj} \circ f)$ where $\mathsf{dist}_\tau(m) = \min_{f'} \mathsf{dist}_\tau(f')$ is the optimal distortion of all randomized approval voting rules.*

Proof. As proportional runoff selects the majority winner by a probability of at least $\frac{1}{2}$ (at least $n/2$ voters prefer the majority winner), $\mathrm{E}[\mathsf{SW}(\mathsf{prop} \circ f)] \geqslant \mathrm{E}[\frac{1}{2} \mathsf{SW}(\mathsf{maj} \circ f)]$, which yields the second inequality.

The first inequality follows from the fact that $\mathrm{E}[\mathsf{SW}(\mathsf{prop} \circ f)] \leqslant \mathrm{E}_{(c_1,c_2) \sim f}[\mathsf{SW}(c_1) + \mathsf{SW}(c_2)] = 2 \cdot \mathrm{E}_{c \sim f}[\mathsf{SW}(c)]$, where the last term is an expectation with marginal probabilities that c appears in the pair returned by f. This is similar to drawing two candidates from a single-stage voting rule. Thus, any lower bound on distortion of single-stage approval voting, holds for approval voting with runoff by an additional factor of $1/2$. □

By the above observation the prior bounds on the distortion of approval voting with or without majority runoff, we conclude the following bounds for the setting with proportional runoff.

Corollary 7. *In approval voting with proportional runoff, $\tau = \frac{1}{m}$, there exists a randomized rule that achieves distortion of $O(m)$. Furthermore, any randomized rule incurs a distortion of*

- *$\Omega(m)$ for $\tau \leqslant \frac{1}{m}$,*
- *and $\Omega(\frac{1}{\tau} + m\tau)$ for $\tau \geqslant \frac{1}{m}$.*

Proof. By Proposition 1, the upper bound follows from Theorem 10 and the lower bound follows from Theorem 5. □

So far, we have shown that distortion is $\Omega(\sqrt{m})$. In Theorem 13, we show a better lower bound for range of values of τ specifically for the setting with proportional runoff. Due to space limits, we defer the proof to Appendix A.6.

Theorem 13. *In approval voting with proportional runoff, for $\tau \in [\frac{3}{m-1}, 1]$, any randomized rule incurs a distortion of at least $\Omega\big(\min\{m, (m\tau)^{3/2}\}\big)$.*

For $\tau \geqslant \frac{1}{\sqrt{m}}$, Corollary 7 shows a lower bound of $\Omega(m\tau)$ which is outperformed by $\Omega((m\tau)^{3/2})$ upper bound in Theorem 13. For $\tau \in [\frac{3}{m}, \frac{1}{\sqrt{m}}]$, Corollary 7 shows a lower bound of $\Omega(\frac{1}{\tau})$ while Theorem 13 shows a $\Omega((m\tau)^{3/2})$. The former is a decreasing function as τ increases and the latter is an increasing function. By combining the two, we get a lower bound of $\max\{\frac{1}{\tau}, (m\tau)^{3/2}\}$ which is minimized at $\frac{1}{\tau} = (m\tau)^{3/2}$, i.e., $\tau = \big(\frac{1}{m}\big)^{3/5}$. Thus, we get a lower bound of $\Omega(m^{0.6})$ for distortion in this setting.

Corollary 8. *In approval voting with proportional runoff, the optimal distortion value for randomized rules is in the range of $[O(m), \Omega(m^{0.6})]$. The upper bound of $O(m)$ is achievable at $\tau = \frac{1}{m}$.*

7 Experiments

In this section, we evaluate the empirical performance of several rules in approval voting with or without runoff. We measure the average-case approximation-ratio between the social welfare achieved by these rules and the optimal social welfare over synthetic and real-world datasets.

Rules. For two-stage rules, we investigate pair selection rules introduced by Delemazure et al. [13]. Each rule in this class selects a pair $\{c, c'\}$ of the candidates maximizing $S_{\vec{\pi}}(c) + S_{\vec{\pi}}(c) - \alpha S_{\vec{\pi}}(c, c')$ for some $\alpha \in [0, 1]$ where $S_{\vec{\pi}}(c)$ is the number of the candidates that approve c. This class is named $\alpha - AV$ rules. We use three rules from this class: MAV ($\alpha = 0$), PAV ($\alpha = 1/2$), and CCAV($\alpha = 1$). We also consider the sequential versions of these rules where the first selected candidate is the most approved candidate and the second candidate is selected to maximize the desired objective given a fixed first candidate. The sequential version of MAV is the same as MAV, but for PAV and CCAV we name their sequential version SPAV and SCCAV respectively. We use both each of the rules defined in composition with majority and proportional runoff. We also consider three single stage rules. One deterministic rule: most approved candidate (MAC), and two randomized rules: (1) proportional to approval score (PAS) and (2) with probability 1/2 most approved candidate, with 1/2 uniform (HMHU). In total we have 13 different rules.

7.1 Synthetic Data

Data Generation. We use $n = 200$ voters. For an instance with m candidates, we create an approval profile as follows. We first generate m random permutations over the alternatives. For each voter, draw one of these permutations as her preference ranking, and draw a utility vector i.i.d. from a Dirichlet distribution with m concentration parameters all set to 1, i.e. Dir(1, ..., 1), and assign utilities of the drawn vector according to the preference ranking. We generate the approval votes based on τ. For each $m \in \{5, 10, 15 \ldots, 100\}$ and $\tau \in \{0, 0.01, 0.02, \ldots, 0.1, 0.15, 0.2, \ldots, 1\}$, we report the average welfare ratio achieved by the rules over 1000 generated instances. The error bars show the standard error.

Results. In Figures 1 and 5 in Section 7.1 we observe that the average distortion of $\alpha - AV$ rules has little dependence on whether we choose the two finalists by maximizing the objective globally or sequentially. We can also see that using majority runoff almost always gives us a better welfare ratio. However, only for larger values of τ, there is a considerable gap between the rules with majority runoff and rules with proportional runoff. Furthermore, for small values of τ, a single-stage rule gives us lower social welfare (higher distortion), but for larger values of τ it is better than proportional runoff. However, majority runoff performs better on these instances.

For a specific m the distortion decreases up to some τ and then starts to increase. In Figure 6 in Section 7.1 we can see that for different values of m the minimum distortion happens at different τ. The question is that, what is the optimal value of τ for each m. We answer this question with Figure 2, where you can see that the best τ is asymptotically $2/m$ for all the rules except PSA. We can see the distortion of different rules with this τ in Figure 7.

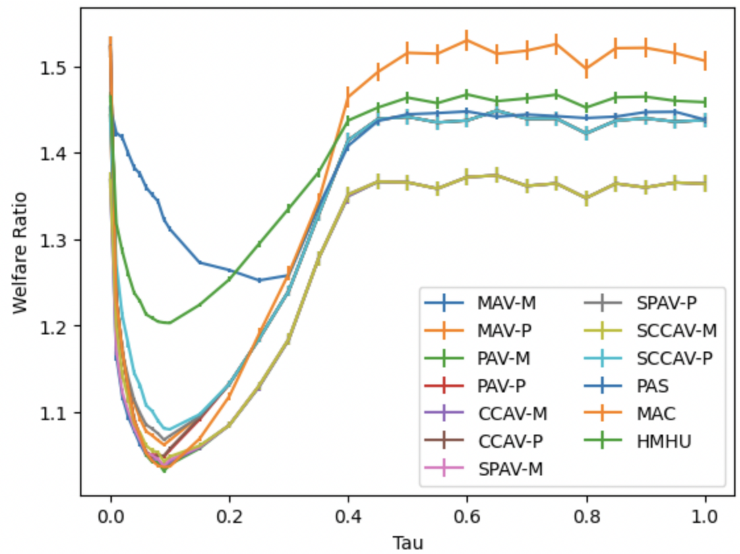

Figure 1: Average distortion of different rules on synthetic data with $m = 20$ candidates.

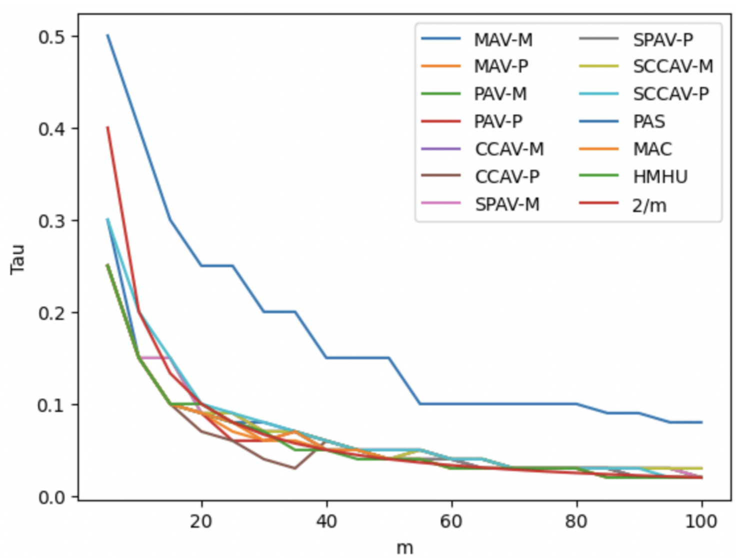

Figure 2: The value of τ that gives the best distortion for different voting rules.

7.2 PrefLib Data

We run our experiment on 134 different real world data from PrefLib [23]. Each instance gives us the number of the candidates and the ranking of the voters over them. For each voter we generate a utility function uniformly drawn from the m-simplex, and change the order of the utilities to make it consistent with the preference ranking. Based on these utilities and with $\tau = 2/m$ we generate the approval votes and run our rules. The welfare ratio of some of the rules on this data for different values of m is presented in Figure 3. We exclude values of m for which less than 3 instances are available in the dataset.

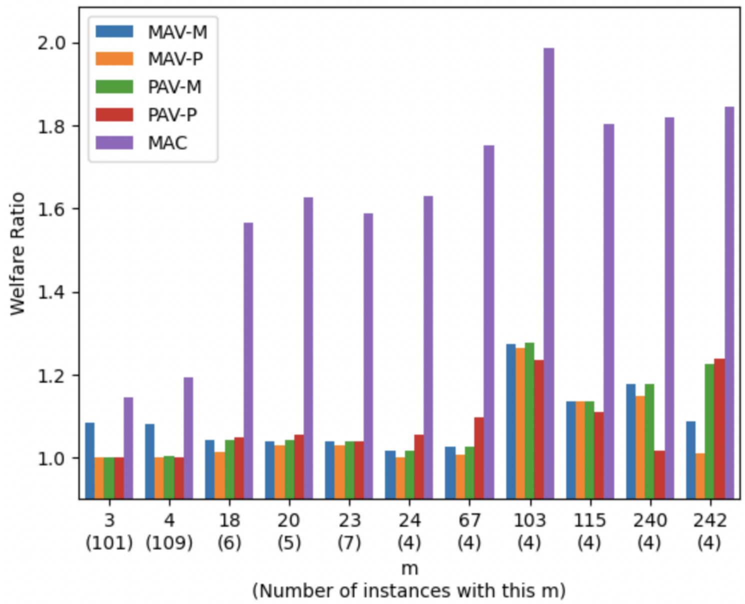

Figure 3: The welfare ratio of different instances from PrefLib SOC data with $\tau = 2/m$.

8 discussion

We studied the class of approval with runoff rules based in the utilitarian distortion framework. We show that, compared to the single-stage setting, distortion increases substantially when a majority runoff is used in the second round. We proposed the randomized proportional runoff system, and showed that for deterministic rules, it keeps distortion as low as it is in the single-stage setting. However, for randomized rules, we leave the optimal distortion value of approval voting with proportional runoff as an open question.

References

[1] 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.

[2] Elliot Anshelevich and John Postl. 2017. Randomized social choice functions under metric preferences. *Journal of Artificial Intelligence Research* 58 (2017), 797–827.

[3] Enriqueta Aragones, Itzhak Gilboa, and Andrew Weiss. 2011. Making statements and approval voting. *Theory and decision* 71, 4 (2011), 461–472.

[4] Haris Aziz. 2020. Justifications of welfare guarantees under normalized utilities. *ACM SIGecom Exchanges* 17, 2 (2020), 71–75.

[5] Haris Aziz. 2020. Justifications of welfare guarantees under normalized utilities. *ACM SIGecom Exchanges* 17, 2 (2020), 71–75.

[6] Gerdus Benade, Swaprava Nath, Ariel D Procaccia, and Nisarg Shah. 2021. Preference elicitation for participatory budgeting. *Management Science* 67, 5 (2021), 2813–2827.

[7] Gerdus Benade, Swaprava Nath, Ariel D Procaccia, and Nisarg Shah. 2021. Preference elicitation for participatory budgeting. *Management Science* 67, 5 (2021), 2813–2827.

[8] Allan Borodin, Daniel Halpern, Mohamad Latifian, and Nisarg Shah. 2022. Distortion in voting with top-t preferences. In *Proc. of 31st IJCAI*. 116–122.

[9] 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.

[10] Steven J Brams and M Remzi Sanver. 2009. Voting systems that combine approval and preference. In *The mathematics of preference, choice and order*. Springer, 215–237.

[11] 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.

[12] Ioannis Caragiannis and Ariel D Procaccia. 2011. Voting almost maximizes social welfare despite limited communication. *Artificial Intelligence* 175, 9-10 (2011), 1655–1671.

[13] Théo Delemazure, Jérôme Lang, Jean-François Laslier, and Remzi Sanver. 2022. Approval with Runoff. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*. International Joint Conferences on Artificial Intelligence Organization, 230–236.

[14] Soroush Ebadian, Anson Kahng, Dominik Peters, and Nisarg Shah. 2022. Optimized distortion and proportional fairness in voting. In *Proceedings of the 23rd ACM Conference on Economics and Computation*. Association for Computing Machinery, New York, NY, USA, 563–600.

[15] Brandon Fain, Ashish Goel, Kamesh Munagala, and Nina Prabhu. 2019. Random dictators with a random referee: Constant sample complexity mechanisms for social choice. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 1893–1900.

[16] Aris Filos-Ratsikas, Evi Micha, and Alexandros A Voudouris. 2020. The distortion of distributed voting. *Artificial Intelligence* 286 (2020), 103343.

[17] Aris Filos-Ratsikas and Alexandros A Voudouris. 2021. Approximate mechanism design for distributed facility location. In *International Symposium on Algorithmic Game Theory*. Springer, 49–63.

[18] Mohammad Ghodsi, Mohamad Latifian, and Masoud Seddighin. 2019. On the distortion value of the elections with abstention. In *Proc. of 33rd AAAI*, Vol. 33. 1981–1988.

[19] Vasilis Gkatzelis, Daniel Halpern, and Nisarg Shah. 2020. Resolving the optimal metric distortion conjecture. In *Proc. of 61st FOCS*. 1427–1438.

[20] D Marc Kilgour, Steven J Brams, and M Remzi Sanver. 2006. How to elect a representative committee using approval balloting. In *Mathematics and democracy*. Springer, 83–95.

[21] Fatih Erdem Kizilkaya and David Kempe. 2022. Plurality Veto: A Simple Voting Rule Achieving Optimal Metric Distortion. In *Proc. of 31st IJCAI*. 349–355.

[22] Martin Lackner and Piotr Skowron. 2018. Approval-Based Multi-Winner Rules and Strategic Voting.. In *IJCAI*. 340–346.

[23] Nicholas Mattei and Toby Walsh. 2013. Preflib: A library for preferences http://www. preflib. org. In *International conference on algorithmic decision theory*. Springer, 259–270.

[24] Grzegorz Pierczyński and Piotr Skowron. 2019. Approval-based elections and distortion of voting rules. *arXiv preprint arXiv:1901.06709* (2019).

[25] Ariel D Procaccia and Jeffrey S Rosenschein. 2006. The distortion of cardinal preferences in voting. In *Proc. of 10th CIA*. 317–331.

[26] M Remzi Sanver. 2010. Approval as an intrinsic part of preference. In *Handbook on approval voting*. Springer, 469–481.

A Missing Proofs

A.1 Proof of Theorem 7

Proof. Consider a τ-approval profile $\vec{\pi}$ in which none of the voters approve any candidates. Without loss of generality, suppose a deterministic rule selects c_1 and c_2 as finalists. Regardless of the outcome of the runoff election, consider a consistent utility profile $\vec{u}$ in which every voter has utility $\frac{1}{m-2}$ for every candidate except c_1 and c_2, and zero utility for c_1 and c_2. Note that this is consistent with $\vec{\pi}$ due to $\tau > \frac{1}{m-2}$. Hence, $\mathrm{SW}(c_1, \vec{u}) = \mathrm{SW}(c_2, \vec{u}) = 0$, meaning that the social welfare of the final winner is zero regardless of the outcome of the runoff election. In contrast, every other candidate have a positive social welfare, yielding unbounded distortion. $\square$

A.2 Proof of Theorem 9

Proof. Similar to Example 1, consider an election in which the top vote of $\frac{n}{2} - 1$ voters is c_1 with utility of 1, while c_1 is the least preferred candidate for the rest with a utility of 0. Hence, $\mathrm{SW}(c_1) = \frac{n}{2} - 1$; and c_1 cannot be the winner of the election as it is the Condorcet loser.

For the rest of the $\frac{n}{2} + 1$ voters, similar to the example in Theorem 3, suppose $\frac{1}{m-1}$ fraction of them have approve $\{c_i, c_{i+1}\}$ for $i \in [2, m-1]$. The remaining $\frac{1}{m-1}$ fraction approve $\{c_m, c_2\}$. These voters have a utility of $1 - \tau$ for their top choice and τ for their second top vote, which can be either c_i or c_{i+1}.

Let c, c' be the finalists determined by the deterministic rule. Suppose $c \neq c_1$, without loss of generality. If we make c the winner of the election, by placing c the second vote of voters who approve c, $\mathrm{SW}(c) \leqslant \frac{n}{m} \cdot \tau$; as a result, distortion is lower bounded by $\frac{n/2-1}{n/m \cdot \tau} \geqslant \frac{m}{2\tau} - 1$. It remains to prove that c can win the election. As c_1 loses to all other candidates including c, we assume that $c' \neq c_1$.

Except for voters who approve c', we can rank c above c' as the given approvals allows for such preference rankings to exist. Therefore, c wins the majority comparison if $2 \cdot \frac{n/2+1}{m-1} < \frac{n}{2}$, which holds for $m > 3$ and $n > 4$. $\square$

A.3 Proof of Lemma 1

Proof. Suppose c is the Condorcet loser in an instance. Let $\mathrm{rank}_i(c) \in [1, m]$ be the rank of c in the preference ranking of voter i. Then, by the unit-sum assumption, we have $u_i(c) \leqslant \frac{1}{\mathrm{rank}_i(c)}$ (otherwise, there are $\mathrm{rank}_i(c)$ candidates worth more than $\frac{1}{\mathrm{rank}_i(c)}$ to voter i which violates the unit-sum assumption). As c loses to all other agents, it must be ranked lower than each candidate by at least $n/2$ voters. Hence, $\sum_i \mathrm{rank}_i(c) \geqslant (m-1) \cdot \frac{n}{2} + n$ (the additional n stems from the fact that the rank functions starts from 1). Then,

$$\mathrm{SW}(i) \leqslant \sum_i \frac{1}{\mathrm{rank}_i(c)} \overset{(1)}{\leqslant} \frac{n}{2} \cdot 1 + \frac{n}{2} \cdot \frac{1}{m},$$

where (1) holds as follows. Due to the convexity of the function $\frac{1}{x}$ and given that $\sum_i \mathrm{rank}_i(c) \geqslant \frac{n}{2}(m-1) + n$, $\sum_i \frac{1}{\mathrm{rank}_i(c)}$ is maximized when some ranks are 1 and the rest are m. That is, it is best to have c the top choice of $\frac{n}{2}$ voters and the bottom choice of the remaining voters, which completes the proof. $\square$

A.4 Proof of Theorem 10

Proof. Consider any instance with utility profile $\vec{u}$ and approval profile $\vec{\pi}$. Let q_1 and q_2 be the number of approvals of c_1 and c_2 respectively. Let c^* be the optimal candidate.

- Case c_1 wins over c_2. By Theorem 2, choosing c_1 deterministically would have bounded the distortion by $2m$. In this case, we choose it with probability of at least $1/2$. Therefore, distortion is bounded by $4m$.

- Case c_1 loses to c_2 and $c^* \neq c_1$. As c_1 loses to c_2, then c_1 is not the top choice of at least $n/2$ voters. By the choice of $\tau = \frac{1}{m}$, every voter must approve at least one candidate. Therefore, there are at least $n/2$ approvals for candidates other than c_1. Therefore, as c_2 is the second most approved candidate, $q_2 \geqslant \frac{n}{2(m-1)}$. Furthermore, by the assumption that $c^* \neq c_1$, c^* has at most q_2 approvals. Hence, $\mathrm{SW}(c^*, \vec{u}) \leqslant q_2 + (n - q_2) \cdot \frac{1}{m}$, and distortion of this rule is bounded by

$$\frac{q_2 + (n - q_2)/m}{q_2/m} \leqslant m + 2m - 1,$$

where the last inequality holds by $q_2 \geqslant \frac{n}{2m}$.

- Case c^* loses to c_2 and $c^* = c_1$. If c_1 wins the majority runoff from any candidate, it will be selected by a probability of at least $\frac{1}{2(m-1)}$, which results in a distortion of at most $2(m-1)$. Otherwise, c_1 is the Condorcet loser. By Lemma 1, $\mathrm{SW}(c_1, \vec{u}) \leqslant \frac{n}{2} + \frac{n}{2m}$. Furthermore, we have selected any other candidate with probability at least $\frac{1}{2(m-1)}$. Hence, $\mathrm{E}[\mathrm{SW}(f, \vec{u})] \geqslant \frac{1}{2m} \cdot (n - \mathrm{SW}(c, \vec{u}))$. Then,

$$\mathsf{dist}_{\frac{1}{m}}(f) \leqslant \frac{\mathrm{SW}(c_1, \vec{u})}{\frac{1}{2m}(n - \mathrm{SW}(c, \vec{u}))} \leqslant \frac{2m}{\frac{n}{\mathrm{SW}(c)} - 1} \leqslant \frac{2m}{3/2 - 1} = 4m,$$

where the last inequality comes from that $\frac{n}{2} + \frac{n}{2m} \leqslant \frac{2n}{3}$ for $m \geqslant 3$. □

A.5 Proof of Theorem 12

Proof. Suppose that for each $i \in [m]$, there are n/m voters who approve candidates c_i, $c_{(i \bmod m)+1}$, and $c_{(i+1 \bmod m)+1}$. Every voter has utility $1 - 2 \cdot \tau$ for their top choice, τ for their second and third choice, and 0 for the rest.

Let c_1 and c_2 be the candidates selected by a deterministic pair-selection rule. Let the $3n/m$ voters approving c_1 or c_2 have utility τ for c_1 and c_2 and utility $1 - \tau$ for their other approved candidate. Pick any candidate $c^* \neq c$, and let the $3n/m$ voters approving c^* have utility $1 - \tau$ for c^* and τ for their other approved candidate.

Under this utility profile $\vec{u}$, $\mathrm{SW}(c_1, \vec{u}) = \mathrm{SW}(c_2, \vec{u}) = \tau \cdot 3n/m$ whereas $\mathrm{SW}(c^*, \vec{u}) = (1 - \tau) \cdot 3n/m$. Hence, the distortion is at least

$$\frac{(1 - \tau) \cdot 3n/m}{2 \cdot \tau \cdot 3n/m} = \frac{1}{2}\left(\frac{1}{\tau} - 1\right).$$

□

A.6 Proof of Theorem 13

First, we prove a technical lemma.

Lemma 2. *Let X be the distribution over $\Delta(C^2)$. There exists a subset $S \subset C$ of size $|S| \leqslant \frac{m}{3}$, such that*

- $\forall c \in S, \Pr[c \in X] \leqslant \frac{6}{m}$,
- *and* $\sum_{c,c' \in S} \Pr[(c, c') \in X] \leqslant 9|S|^2/m^2$.

Proof. Let $T = \{c \mid \Pr[c \in X] \leqslant 6/m\}$. Then, $|T| \geqslant \frac{2m}{3}$, otherwise, there at least $\frac{m}{3}$ candidates with marginal probability of more than $\frac{6}{m}$, which is contradiction as the sum of marginal probabilities for all candidates is 2. We construct the sets S^k for all $k \in [1, \frac{m}{3}]$ iteratively using only candidates in T, described as follows. Start with an empty set S^0. Pick a random member of T and add to S^1. Define $\Pr[S^k] = \sum_{c \in S^k} = \sum_{\Pr[c \in X]}$ as the sum of marginal probabilities of candidates in S^k. If $S^k \subset T$, then the first condition is satisfied and we also have $\Pr[S^i] \leqslant 6k/m$.

Inductively speaking, at the k-th iteration, there exists a candidate $c^k \in T \setminus S^{k-1}$, such that $\sum_{c' \in S^{k-1}} \Pr[(c', c^k)] \leqslant \frac{\Pr[S^{k-1}]}{|T \setminus S^{k-1}|} \leqslant \frac{6(k-1)/m}{m/3} = 18(k-1)/m^2$. Let $S^k = S^{k-1} \cup \{c^k\}$. By such construction, we have $\sum_{c,c' \in S^k} \Pr[(c, c')] \leqslant \sum_{i=1}^{k} \frac{18(i-1)}{m^2} \leqslant \frac{9k^2}{m^2}$, which satisfies the second condition. □

Now, we show the lower bound construction.

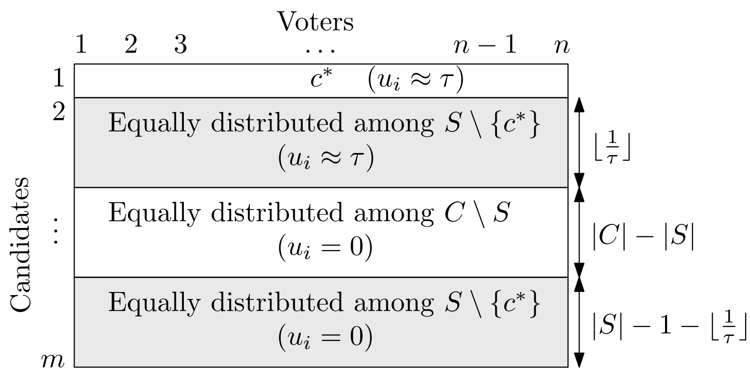

Figure 4: Preference and utility profile of the lower bound example described in Theorem 13.

Proof of Theorem 13. Let f be a randomized pair selection rule. We construct an example for f as follows (depicted in Figure 4). Suppose voters have an empty approval set for threshold τ, and that each voter divides their total utility of one between their $\ell = \lfloor \frac{1}{\tau} \rfloor + 1$ top candidates (which we define in the next paragraph). This way, each voter has a utility of $\frac{1}{\ell} \in [\frac{\tau}{2}, \tau)$ for their ℓ top candidates and 0 for the rest.

Let $k = \sqrt{\frac{2}{3} m(\ell - 1)}$. Select a set S of $k+1$ candidates according to Lemma 2. Pick one candidate from S $c^* \in S$ (the last one added by the iterative construction) and place c^* as the top vote of all voters. This way, c^* would be the social welfare maximizing candidate with $\mathrm{SW}(c^*) = \frac{n}{\ell}$ that is selected by the rule with probability at most $\frac{6}{m}$. Suppose the other $\ell - 1$ top positions of all voters is divided equally between $S \setminus \{c^*\}$. Therefore, for all $c \in S \setminus \{c^*\}$, we have c in $\frac{n(\ell-1)}{k}$ positions worth $\frac{1}{\ell}$, i.e., $\mathrm{SW}(c) = \frac{n(\ell-1)}{k} \cdot \frac{1}{\ell} \leqslant \frac{n}{k}$. Place other candidates $C \setminus S$ in the positions $l+1$ to $m - \ell + 1$, and fill the bottom of the rankings with $S \setminus \{c^*\}$. This way, for each $c' \in C \setminus S$, the proportional runoff between $c \in S \setminus \{c^*\}$ and c' results in a winning chance, for c, of $\Pr[c \in \mathsf{prop}(c, c')] = \frac{1}{n} \cdot \frac{n(\ell-1)}{k} = \frac{\ell-1}{k}$.

Now, we show an upper bound on the social welfare obtained f. The rule selects c^* with probability of at most $\frac{6}{m}$ (true for all in S), which contributes $\frac{6}{m} \cdot \frac{n}{\ell}$ to the obtained welfare. For each $c \in S \setminus \{c^*\}$, the probability of c winning the election is

$$\begin{aligned}\Pr[c \in \mathsf{prop} \circ f] &= \sum_{c' \in C} \Pr[(c, c') \in f] \cdot \Pr[c \in \mathsf{prop}(c, c')] \\ &\leqslant \Big(\sum_{c' \in S \setminus \{c^*\}} \Pr[(c, c') \in f] \Big) + \frac{6}{m} \cdot \frac{\ell - 1}{k}.\end{aligned}$$

By Lemma 2, we have that $\sum_{c,c' \in S \setminus \{c^*\}} \leqslant \frac{9k^2}{m^2}$. Thus,

$$\sum_{c \in S \setminus \{c^*\}} \Pr[c \in \mathsf{prop} \circ f] \leqslant \frac{9k^2}{m^2} + k \cdot \frac{6(\ell - 1)}{mk}.$$

Therefore,

$$\begin{aligned}
\mathrm{E}[\mathrm{SW}(\mathsf{prop}\circ f)] &\leqslant \frac{6}{m}\cdot\frac{n}{\ell} + \sum_{c\in S\setminus\{c^*\}} \Pr[c\in \mathsf{prop}\circ f]\cdot\frac{n}{k} \\
&\leqslant \frac{6}{m}\cdot\frac{n}{\ell} + \frac{9nk}{m^2} + \frac{6n(\ell-1)}{mk} \\
&\leqslant \frac{6}{m}\cdot\frac{n}{\ell} + \frac{3n}{m}\cdot 6\sqrt{\frac{2(\ell-1)}{3m}}.
\end{aligned}$$

Putting all together, distortion of f is lower bounded by

$$\begin{aligned}
\frac{n/\ell}{\frac{6}{m}\cdot\frac{n}{\ell}+\frac{3n}{m}\cdot 6\sqrt{\frac{2(\ell-1)}{3m}}} &\geqslant \frac{1}{2}\min\{\frac{m}{6},\frac{m^{3/2}}{15\cdot\ell^{3/2}}\} \\
&\geqslant \Omega\left(\min\{m,(m\tau)^{3/2}\}\right).
\end{aligned}$$

□

B Missing Experimental Results

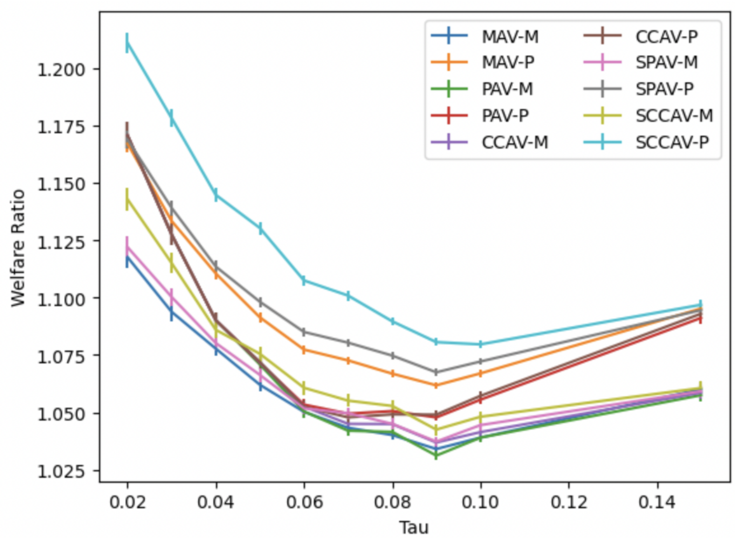

Figure 5: Close look at the average distortion of α-AV rules on synthetic data with $m = 20$ candidates, and small τ.

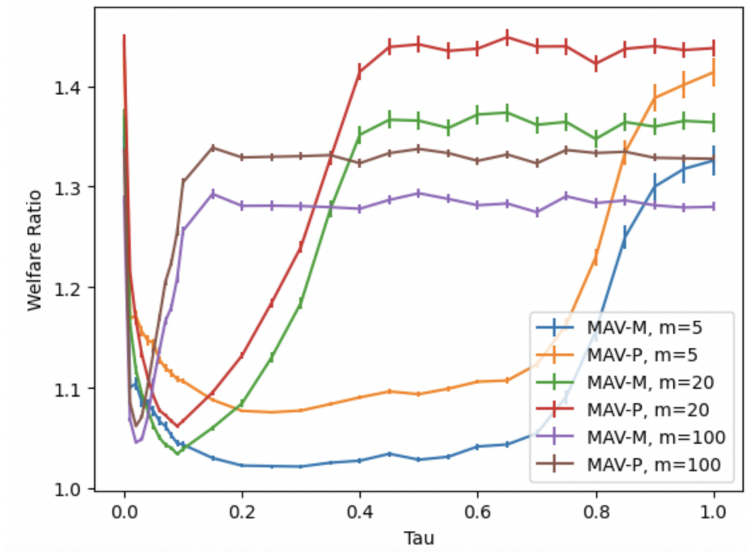

Figure 6: Average distortion of different α-AV rules for different number of candidates.

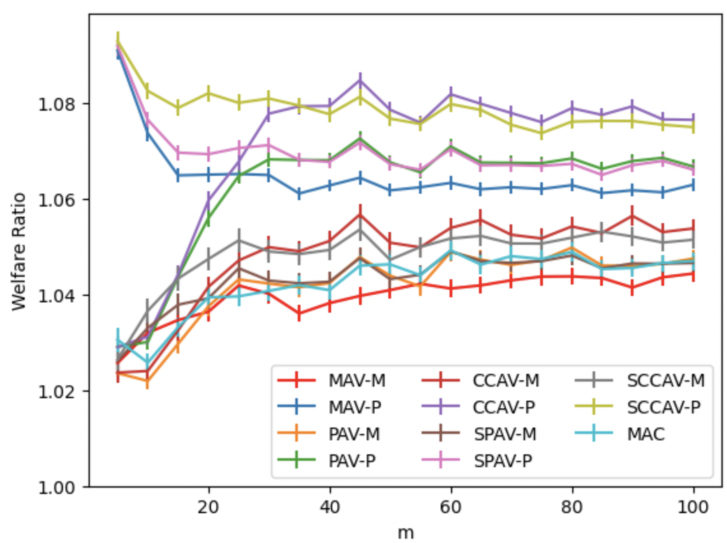

Figure 7: The welfare ratio of different rules with $\tau = 2/m$.