Voters with Stakes Can Ward Off Bad Candidates

BAILEY FLANIGAN, ARIEL PROCACCIA, and SVEN WANG

The *distortion* of a voting rule measures how bad (in terms of social welfare) the winning candidate or alternative can be due to limited information about utilities. Without any further assumptions, unbounded distortion is inevitable. We argue, however, that it is reasonable to assume that some information about the *stakes* of voters is known, where someone with high stakes stands to gain or lose a lot, depending on the outcome. We develop a general framework that captures this idea through the notion of *stakes-proportionality*, whereby voters are reweighted according to their stakes. Using this framework, we derive tight bounds on achievable distortion under stakes-proportionality, when perfect information about stakes is available. When stakes can only be estimated, we show that our distortion bounds hold approximately, implying that even a coarse estimate of stakes can bring about a large reduction in distortion. Finally, for the setting where stakes are completely unknown, we develop a proof-of-concept mechanism that ties together multiple elections so that the behavior of voters reveals their stakes, leading to bounded distortion in each election.

Manuscript submitted for review to the 25th ACM Conference on Economics & Computation (EC'24).

1 INTRODUCTION

In the standard model of voting, voters express their *ordinal* preferences by ranking a set of alternatives. It is reasonable to assume, however, that there *exist* cardinal utilities that voters associate with alternatives, and a voter's ranking is consistent with those latent utilities. From this viewpoint, a natural goal for a voting rule — which aggregates the given rankings — would be to select a good alternative in terms of utilitarian social welfare (the sum of utilities), despite having access only to ordinal information. The notion of *distortion* [Procaccia and Rosenschein, 2006] measures how far a given voting rule is from achieving this goal. It is the worst-case ratio between the social welfare of the optimal alternative and the social welfare of the alternative selected by the rule, where the worst case is taken over utilities.

Without any additional assumptions, any deterministic voting rule must have unbounded distortion. Intuitively, even in an election with two voters, one preferring a to b and the other preferring b to a, it could be the case that one voter has arbitrarily high utility for their preferred alternative whereas the other has low utility for both alternatives, but it is impossible to differentiate between a and b based on the ordinal information alone. To circumvent this obstacle, the rich literature on distortion has explored several approaches, one of which is to assume access to limited cardinal information (see Section 1.2 for details). Our work is inspired by this approach; we are interested in a specific and very natural type of cardinal information: *stakes*.

To motivate the idea of stakes, let us consider a concrete example. Cambridge, Massachusetts, like many US cities, seeks public input on the questions surrounding affordable housing. In past years, the city has proposed to permit building taller buildings — a move that would increase affordable housing, but would change the skyline of Cambridge and potentially cast shadows over existing homes. City Councilor Burhan Azeem, who proposed the amendments, said "I understand that tall buildings are something that people are sensitive to, but this comes down to which should we care more about. How tall a building is? Or the people who don't have stable housing?" [Brinker, 2023]. Councilor Azeem is pointing out a contrast of *stakes*: while both homeowners and those who don't have stable housing are affected by the decision of whether to permit taller buildings, the latter group is *far more affected*.

Now, suppose the decision is put to public referendum; placing this example in the utility model above (using numbers to illustrate the intuition), suppose that 10% of residents are renters who are at risk of losing their home due to rising rental prices, and 90% are homeowners who are worried that affordable housing would lead to a decrease in the value of their property. Let the former group have utility 100 for accepting the proposal and utility 0 for rejecting it; let the latter group have utility 1 for rejecting the proposal and utility 0 for accepting it. Based on their utilities, 90% of residents will vote to reject the proposal, and any majority-consistent voting rule must confirm their choice. This is a severely suboptimal outcome, as the social welfare of accepting the proposal is more than 10 times that of rejecting it.

Notice that, in the foregoing example, the chosen numbers do not matter much: the key problem is that a minority of residents have disproportionately high stakes, but they lack the voting power to sway the election. By contrast, the majority stands to gain little, so when they get their way, little value is generated for the population. It is not hard to think of salient examples of real political decisions with this property: consider masking requirements and immunocompromised populations, accessibility features (e.g., wheelchair ramps) and those who rely on them, or the design of public transit and those who cannot afford private transportation. Although the impossibility that all deterministic rules have unbounded distortion is often thought of as a theoretical one, these examples make the issue seem practically pressing: given that people *are* likely to have disparate stakes in real issues — and sometimes minority groups have much higher stakes than the majority — these examples suggest that such welfare loss can occur in *real elections*.

Motivated by this problem, we pursue bounded distortion by assuming that voters' stakes are known, at least approximately, to the voting rule. Unlike a significant branch of the distortion literature that assumes voters' utilities are normalized, we will allow voters' utilities to be arbitrary and unknown to ensure that our model captures the problematic examples above. One may wonder whether it is plausible to know stakes information but not precise utilities. To see why knowing stakes is far easier, consider the affordable housing example: conceptually, stakes captured *how affected each voter is relative to other voters*. Stakes information, then, is just a single number measuring the extent to which a voter can gain or lose depending on the decision outcome, *relative to other voters in the election*. As the number of alternatives in an election grows, the gap in difficulty widens between the doing a coarse-grained assessment of how relatively affected a voter is (which in the examples above is not hard to do) versus understanding individuals' fine-grained preferences over all alternatives. As we discuss below, our results permit having just approximate stakes information; going further, we then initiate a study of how voters' stakes can be revealed in their behavior, potentially permitting our positive results to hold even in cases where legitimate estimates of voters' stakes are unavailable.

1.1 Approach and contributions

How should we measure stakes? (Section 2). First, we must embed a model of stakes in the standard model of voting with latent utilities. The affordable housing example illustrates that, intuitively, a voter's stakes are captured in their *utility vector* — that is, their utilities across alternatives. In that example, it seems natural to measure voters' stakes as the difference between their utilities (so the respective stakes of renters and homeowners would be 100 and 1). However, how to measure stakes becomes less obvious when $m > 2$. (Consider the utility vectors $(1, 1, 0)$ and $(1, 0, 0)$. Which reflects higher stakes?) We thus define and study general *stakes functions* s: any mapping from a voter's utility vector to a scalar measure of their stakes in the election.

What can we do with *perfect* stakes information? (Section 3). In the affordable housing example, an intuitive idea for addressing high distortion would have been to re-weight votes according to voters' stakes. Taking this approach, we characterize the distortion possible, across stakes functions s, by any deterministic or randomized rule when votes are re-weighted in a stakes-proportional way.

Deterministic rules. Here we find that knowing stakes information — even according to a surprisingly simple stakes function — can drop the distortion from ∞ to the number of alternatives m. We first prove a lower bound showing that no deterministic voting rule, when reweighted by *any stakes function* s according to *any reweighting scheme*, can achieve distortion lower than m (Theorem 3.1). We then prove a general upper bound on the distortion of *any* deterministic voting rule, when votes are reweighted proportionally via *any* stakes function s (Theorem 3.4). We use this bound to identify a stakes function, voting rule pair that matches our lower bound: we show that the rule Plurality achieves optimal distortion m when reweighted according to the stakes as measured by *maximum utility* or the *difference between maximum and minimum utilities* (Proposition 3.7). We henceforth refer to these stakes functions as $s = \mathsf{max}$ and $s = \mathsf{range}$, respectively. This bound further implies proportional re-weighting is sufficient to achieve optimality. Beyond Plurality, we surprisingly find that among deterministic voting rules, most are not helped by knowledge of stakes — a result which has both positive and negative interpretations.

Randomized rules (plus, an independently interesting lemma). We repeat this analysis for randomized rules, showing that if a stakes-reweighted voting rule is permitted to be randomized, the distortion can be as low as $O(\sqrt{m})$. We show that the rule Stable Lottery [Ebadian et al., 2022] achieves $O(\sqrt{m})$ distortion when given stakes measured by $s = \mathsf{max}$, $s = \mathsf{range}$, or $s = \mathsf{sum}$ (the *sum* of a voter's utilities) (Theorem 3.12). The lemma used to prove this upper bound may be of independent interest, as it shows a surprising and technically useful connection between our setting and the popular setting of distortion assuming *unit-normalized utilities* (i.e., where voters' utilities are assumed to sum to 1). We observe, first, that assuming voters' utilities sum to 1 is akin to assuming they have *identical stakes*, as measured by the stakes function $s = \mathsf{sum}$. Then, we show that this assumption is *equivalent*, from a distortion perspective, to permitting arbitrary utilities but reweighting votes by stakes measured by $s = \mathsf{sum}$. In fact, we show this equivalence holds for *any 1-homogeneous stakes function* (Lemma 3.13). This result establishes a formal link — and permits the transfer of bounds — between two key assumption classes in the distortion literature: unit-normalized utilities and queries of cardinal information (our setting). Finally, we conclude our analysis with a lower bound proving that Stable Lottery's distortion is within a $\log m$ factor of optimal across randomized rules and stakes functions (Theorem 3.11).

What if we just have stakes information in *orders of magnitude*? (Section 4). In the motivating examples above, we can identify one or more population groups who are substantially disproportionately affected by the decision. However, assigning any single number to the *extent* to which their stakes are higher is difficult; in reality, people's stakes may be observable (or even fundamentally measurable) only at the resolution of orders of magnitude. Fortunately, in Section 4 we find that even very approximate stakes information can help: we give an instance-wise upper bound showing that for any 1-homogeneous stakes function s (a natural class encompassing all s discussed so far), if our stakes information is δ-approximately correct (and adversarially-designed otherwise), the distortion of any rule f scales simply by δ (Theorem 4.1). Given that the distortion of deterministic voting is otherwise unbounded, this result means that even just coarsely accounting for voters' differing stakes can offer substantial improvements in distortion.

What if we don't have *any stakes information at all?* (Section 5). In some cases, we may not even be able to guess voters' stakes at the resolution of orders of magnitude — or, even more likely, one may encounter cases where it is not possible to make a publicly acceptable case that a certain group should receive disproportionate representation in a given democratic decision. However, in many such cases, it is still desirable to account for stakes; this is an extremely sticky problem, which merits more exploration than we can do in a single paper. We present an initial study of this problem in Section 5, where we propose and explore a class of mechanisms that aims to *reveal voters' information in their behavior* and, in the process, accounts for stakes automatically. This class of mechanisms, called *multi-issue mechanisms*, is based on a simple idea: it requires voters to decide how to allocate a total allotment of voting power across elections. In forcing voters to make such trade-offs, this class of mechanisms exploits a special property of stakes information: unlike other kinds of cardinal utility information, stakes-information is *action-relevant*, describing, roughly, how much a voter may care about the outcome of a given election.

After defining this class of mechanisms, we perform a proof-of-concept analysis of a mechanism in this class. This example illustrates mathematically how voters' stakes can be revealed in their behavior, and how our results from the previous sections can be applied to analyze such a mechanism. Our distortion analysis of this mechanism shows that although any election run individually within the mechanism could have unbounded distortion, the distortion of each election in our mechanism is at most δm^2, where δ corresponds to the extent to which voters have the same *total stakes* across the issues we place on the multi-issue ballot (Theorem 5.4).

1.2 Related work

The literature on distortion is quite rich; for an overview, we refer the reader to the survey by Anshelevich et al. [2021]. At a high level, there are at least three avenues to achieving meaningful bounds on distortion. The first (and by far the most common) is to restrict the utilities, e.g., by assuming that voters have the same sum of utilities [Boutilier et al., 2015, Procaccia and Rosenschein, 2006], or by assuming that the utilities are induced by an underlying metric space [Anshelevich et al., 2018]. The second is to consider *public spirit*, in the sense that voters seek to optimize social welfare in addition to their own utilities [Flanigan et al., 2023]. And the third, which is most relevant to our work, is to assume the availability of limited cardinal information.

In the context of the last approach, the work of Amanatidis et al. [2021] is most closely related to ours. They study deterministic voting rules with access to one of two kinds of queries: *value* queries, where the voting mechanism can directly ask agents about any one of their utilities; and *comparison* queries, where the voting mechanism can ask agents: "for alternatives a and b, is your utility for a at least τ times your utility for b?" Among other results, they prove that constant distortion is achievable using $O(\log^2 m)$ queries per voter, where m is the number of alternatives. To achieve their upper bounds, they construct an approximate utility profile via the queries and maximize social welfare with respect to these estimated utilities. This is conceptually very different from our approach of employing common voting rules and accounting for stakes through stakes-proportionality. Nevertheless, there are a few technical connections between the work of Amanatidis et al. [2021] and ours, which comment on below.

Another (more distantly) related paper in the same vein is that of Abramowitz et al. [2019]. Their main results pertain to a setting where answers to the above comparisons queries are given for every pair of alternatives, either with respect to a single fixed threshold τ or multiple fixed thresholds; their distortion bounds are parameterized by these thresholds. They additionally assume an underlying metric space and therefore their results are technically incomparable to ours.

Although standard social choice mechanisms do not account for stakes,[1] the concept has been conceived of in multiple disciplines. From the social sciences, there is the philosophical notion of *proportionality* — the idea that "power should be distributed in proportion to people's stakes in the decision under consideration" [Brighouse and Fleurbaey, 2010]. In the context of binary decisions, Azrieli [2018] explores the welfare cost of treating agents symmetrically when they have different stakes, and the results of Fleurbaey [2008] suggest that accounting for stakes can help: they show that in elections over two alternatives, reweighting each voter's vote by the difference of their utilities — a measurement of their "stakes" — increases the welfare of majority voting. Our work can be seen as generalizing the latter analysis substantially, permitting m alternatives, any stakes function s, any stakes-reweighting scheme, and any voting rule.

The idea that those with higher stakes in a decision should have greater political influence arises in many other theories in political science as well, such as the *principle of affected interests* [Fung, 2013]; the concept of *empowered inclusion* [Beauvais and Warren, 2019, Warren, 2017]; and the concept of *precarity* [Näsström and Kalm, 2015], which describes the idea that socially or economically vulnerable populations may be more affected by political issues due to their inability to adjust to decisions that are sub-optimal for them. Our work can be seen as a technical companion to this literature in three ways: it (1) offers a formal framework for modeling stakes, (2) explores the impact of accounting for stakes as proposed by these theories, and (3) identifies a class of mechanisms which can be explored further as a method for accounting for stakes.

[1] We know of two social choice papers that use the term "stakes" [Andersen et al., 2014, Kahng and Kehne, 2022]; both use the term differently and explore unrelated questions.

Our mechanism design approach requires assumptions about how voters will engage with the mechanism, which can be informed by existing social science research on how voter behavior depends on preference intensity [Downs, 1957, Feddersen and Pesendorfer, 1996]. Moreover, our proposed class of mechanisms relates to (but is not encompassed by) two existing voting mechanisms, *quadratic voting* [Posner and Weyl, 2015] and *storable votes* [Casella, 2005]. We defer a detailed discussion of these connections to Sections 5 and 6.1, where we can establish them more concretely in comparison to the mechanisms we study.

2 MODEL

We introduce the model in two parts. Section 2.1 establishes the standard voting model; Section 2.2 embeds our model of voters' stakes within it. Throughout the paper, we use the shorthand $\mathbf{1}_\ell\mathbf{0}_{\ell'}$ to mean a vector containing ℓ ones followed by a string of ℓ' zeros. We let $\mathbb{I}(\cdot)$ be the indicator function.

2.1 The voting model

In an election, there are n voters and m *alternatives*. We let voters $i \in [n]$ and alternatives $a \in [m]$ have some fixed numbering. Voters' underlying preferences over $[m]$ are modeled with *utilities*: each voter i has a utility $u_i(a) \in \mathbb{R}_{\geq 0}$ for each alternative $a \in [m]$. Let $\mathbf{u}_i = (u_i(a)|a \in [m])$ be i's *utility vector*, and let $\mathbf{u}$ be a generic utility vector. We summarize all voters' utilities in a *utility matrix* $U \in \mathbb{R}^{n\times m}_{\geq 0}$.

The voting process. Each voter i expresses their preferences via a complete *ranking* over (i.e., permutation of) $[m]$. Letting S_m be the set of all permutations of $[m]$, a generic ranking is $\pi \in S_m$. We use $a \succ_\pi a'$ to denote that a precedes a' in π, reflecting that a is preferred over a'. Abusing notation slightly, we let $\pi(j)$ denote the alternative ranked in the j-th position in ranking π. When voter i "votes", they submit a ranking π_i. This ranking is determined by $\mathbf{u}_i$: i ranks alternatives in decreasing order of their utilities, so that $u_i(a) > u_i(a') \implies a \succ_{\pi_i} a'$ for all $a, a' \in [m]$.[2]

A collection of n voters' rankings is called a preference *profile* $\boldsymbol{\pi}$. As in prior work such as that of Xia [2020], instead of working with profiles $\boldsymbol{\pi}$, we will work with *histograms*, which summarize collections of rankings by their frequencies. A generic preference *histogram* is a vector indexed by rankings, $\mathbf{h} = (h_\pi|\pi \in S_m)$, where $h_\pi \in [0,1]$ is the fraction of rankings in a given collection equal to π. As such, $\|\mathbf{h}\|_1 = 1$. The space of all possible preference histograms is thus the simplex of all valid distributions over S_m, which we call $\Delta(S_m) := \left\{\mathbf{h} \in [0,1]^{S_m} : \sum_{\pi\in S_m} h_\pi = 1\right\}$. Connecting profiles and histograms, we say $\boldsymbol{\pi}$ is *consistent* with $\mathbf{h}$ if each $\pi \in S_m$ appears in $\boldsymbol{\pi}$ exactly $n \cdot h_\pi$ times. Let $\Pi^{\mathbf{h}}$ be the set of all profiles consistent with a histogram $\mathbf{h}$. Note that $\Pi^{\mathbf{h}}$ is non-empty iff all entries of $\mathbf{h}$ are rational.

Since voters' rankings are fully implied by U, we let U constitute an instance. We denote the histogram implied by U as $\mathsf{hist}(U)$, whose π-th entry is given by

$$\mathsf{hist}_\pi(U) := {}^1\!/\!_n \sum_{i\in[n]} \mathbb{I}\{\pi_i = \pi\}, \quad \text{for all } \pi \in S_m.$$

Voting rules. Let $\Delta([m])$ denote the set of all probability distributions over the alternatives $[m]$. Then, a *voting rule* is a function $f : \Delta(S_m) \to \Delta([m])$ that maps a preference histogram to a distribution over winning alternatives.[3] We refer to this class of functions as *randomized* rules to distinguish them from their sub-class, *deterministic* voting rules, which map a histogram to a distribution with singleton support. Among deterministic rules, the only specific rule we study is

[2] For simplicity of our lower bounds, we will assume worst-case rankings when $u_i(a) = u_i(a')$; one could instead tie-break explicitly by perturbing the utilities by arbitrarily small amounts.

[3] Histograms are inherently anonymized, so we study only anonymous voting rules (encompassing all common voting rules).

PLURALITY, whose winner is the alternative that is ranked first by the most voters. Among randomized rules, we consider the STABLE LOTTERY rule [Ebadian et al., 2022], which draws a winner either at random or from a *stable lottery* — a distribution over a subset of $[m]$ that is preferred by voters to other such subsets. We will not apply this rule's precise definition, so we defer it to Appendix B.9.

Distortion. Let an alternative a's utilitarian *social welfare* be $\mathsf{sw}(a, U) := \sum_{i\in[n]} u_i(a)$. We benchmark the social welfare of the winner against that of $a^* := \arg\max_{a\in[m]} \mathsf{sw}(a, U)$, the highest-welfare alternative. For any rule f, the value of this approximation ratio in instance U is called the *instance-specific distortion*, defined as

$$\mathsf{dist}_U(f) := \frac{\mathsf{sw}(a^*, U)}{\mathbb{E}[\mathsf{sw}(f(\mathsf{hist}(U)), U)]},$$

where the expectation is over the draw of the winner from the distribution $f(\mathsf{hist}(U))$. As is standard, we evaluate f via its overall *distortion*, $\mathsf{dist}(f)$, which is the worst-case approximation ratio over all possible instances U:

$$\mathsf{dist}(f) := \sup_{n\geq 1} \sup_U \ \mathsf{dist}_U(f).$$

The supremum over n is just to more conveniently deal with the fact that in worst-case instances, n must be large enough relative to m to realize utility matrices with m-dependent fractional compositions. We consider the distortion to be a function of m, as is standard in the literature.

2.2 A stakes framework within the voting model

Measuring stakes via *stakes functions*. A *stakes function* is any function $s : \mathbb{R}^m_{\geq 0} \to \mathbb{R}$ that maps a utility vector to a scalar measure of the stake it reflects. Intuitively, a voter's stake should depend on the relative magnitudes of their utilities, but not *which* alternatives they prefer; we thus restrict to functions s which are *permutation invariant*. For example, utility vectors $(0, 1)$ and $(1, 0)$ reflect the same stake. We often apply this invariance to evaluate voters' stakes on a *sorted* version of their utility vector.

In some results, we restrict our consideration to stakes functions that are 1-*homogeneous*, i.e., for all scalars α, $s(\alpha\mathbf{u}) = \alpha s(\mathbf{u})$. This applies to $\alpha = 0$, implying that for all 1-homogeneous s, $s(\mathbf{0}) = 0$. This restriction on s is natural in that it makes our notion of *accounting for stakes*, formalized below, invariant to rescaling U. Although many of our results apply for generic stakes functions, three in particular will come up frequently, so we define shorthand for them:

$$\mathsf{range}(\mathbf{u}) := \max_a u(a) - \min_a u(a), \qquad \mathsf{max}(\mathbf{u}) := \max_a u(a), \qquad \mathsf{sum}(\mathbf{u}) := \textstyle\sum_a u(a)$$

Stakes-reweighting of votes. Colloquially, we say that a voting process "accounts for stakes" if it grants voters representation to an extent that depends on their relative stakes. We can think of this as a form of stakes-dependent reweighting: instead of voter i's ranking contributing to the π_i-th entry of the histogram with weight $1/n$, its contribution is additionally weighted by some function of $s(\mathbf{u}_i)$. We can also think of this as *recomposing* the electorate, by effectively duplicating voters in proportion to some function of $s(\mathbf{u}_i)$. For convenience of intuition and notation, we adopt the second interpretation. Formally, let $r : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ be a generic *recomposition function*. Then, the (r, s)-recomposed histogram arising from U has π-th entry

$$\mathsf{hist}^{r\circ s}_\pi(U) = \frac{\sum_{i\in[n]} r(s(\mathbf{u}_i)) \cdot \mathbb{I}(\pi_i = \pi)}{\sum_{i\in[n]} r(s(\mathbf{u}_i))} \quad \forall \pi \in S_m,$$

representing the fraction of voters in the (r, s)-recomposed electorate with ranking π. In this recomposed histogram, each voter i's ranking is represented with weight $r(s(\mathbf{u}_i))/\sum_{i\in[n]} r(s(\mathbf{u}_i))$.

***Proportional* stakes-reweighting of votes.** In this paper, we will focus on perhaps the simplest reweighting scheme: *stakes-proportionality*, where r is the identity function I. In the *s-proportional* histogram $\mathsf{hist}^{I\circ s}(U)$, voters' votes are reweighted *in proportion* to their stakes, i.e., by $s(\mathbf{u}_i)/\sum_{i\in[n]} s(\mathbf{u}_i)$. For notational simplicity, we will shorten the name of $\mathsf{hist}^{I\circ s}$ to hist^s henceforth. We will use *stakes proportionality* (or s-proportionality) to refer to the condition under which votes are reweighted in this way.

Distortion under stakes-proportionality. In a given instance, reweighting an electorate to be stakes-proportionate will potentially change the distortion. We define the *s-distortion* of f as its distortion in $\mathsf{hist}^s(U)$, the s-proportional electorate arising from U:

$$\mathsf{dist}^s_U(f) := \frac{\max_{a\in[m]} \mathsf{sw}(a, U)}{\mathbb{E}[\mathsf{sw}(f(\mathsf{hist}^s(U)), U)]}, \text{ and } \mathsf{dist}^s(f) := \sup_{n\geq 1} \sup_{U} \ \mathsf{dist}^s_U(f).$$

3 WHAT IF WE HAVE *PERFECT* STAKES INFORMATION?

3.1 Deterministic voting rules

We begin by analyzing deterministic voting rules, which without stakes information have infinite distortion (the proof is essentially the housing example in Section 1; for a formal proof, see Appendix B.1). Motivated by this impossibility, we will study what s-distortion is possible across deterministic rules and stakes functions.

3.1.1 Lower bound for all f, s. Our first result is a lower bound that shows that the best possible s-distortion achievable by any deterministic rule f, given stakes information according to any stakes function s, is at least $m-1$.

THEOREM 3.1 (**LOWER BOUND**). *For all s and deterministic f,*

$$\mathsf{dist}^{\mathsf{s}}(f) \geq m - 1.$$

PROOF SKETCH. Our approach is to define two instances, U and U', and show that all deterministic rules f must have at least $m-1$ distortion in one of these two instances. To construct U, U', first set aside one alternative a'; let the remaining alternatives be $a_1 \dots a_{m-1}$. For any $\ell \in [m-1]$, define $A_\ell = \{a_j | j \in [m] \setminus \{\ell\}\}$. When we write A_ℓ in a ranking, it represents a ranking over all the alternatives within it *in increasing order of index.* Now, we define voters' rankings $\boldsymbol{\pi}$ and two possible underlying utilities U and U'. Divide voters in into $m-1$ groups, and consider a voter i in group ℓ. Let them rank alternatives as $\pi_i = a_\ell \succ a' \succ A_\ell$. Their underlying utility vectors as given by U and U', called $\mathbf{u}_i$ and $\mathbf{u}'_i$, are defined below:

alternative:	a_ℓ	$\succ$	a'	$\succ$	A_ℓ
$\mathbf{u}_i$ for $i \in$ group ℓ:	1		1		$0\dots 0$
$\mathbf{u}'_i$ for $i \in$ group ℓ:	1		0		$0\dots 0$

With this construction, the proof follows from three observations. *Observation 1:* First, because both U and U' result in the same ranking for each voter i, $\mathsf{hist}(U) \equiv \mathsf{hist}(U')$, and these two underlying utility matrices are indistinguishable to any voting rule. *Observation 2:* Within each utility matrix, all voters have the same ordered utility vector, and thus have the same stakes; formally, $\mathsf{hist}^s(U) \equiv \mathsf{hist}(U)$ and $\mathsf{hist}^s(U') \equiv \mathsf{hist}(U')$. *Observation 3:* In both U and U', the social welfare of any a_ℓ is equal to $n/(m-1)$; however, in U, a' has high social welfare ($\mathsf{sw}(a', U) = n$) while in U', a' has low social welfare ($\mathsf{sw}(a', U') = 0$).

Now, when a voting rule f receives the profile $\boldsymbol{\pi}$, either $f(\boldsymbol{\pi}) = a'$ or $f(\boldsymbol{\pi}) = a_\ell$ for some $\ell \in [m-1]$. If the former is true, *Observations 1-3* imply that for any s, $\mathsf{dist}^s_{U'}(f) = \frac{n/(m-1)}{0} = \infty$. If the latter is true, *Observations 1-3* imply that for any s, $\mathsf{dist}^s_{U'}(f) = \frac{n}{n/(m-1)} = m-1$. We spell out this argument in full formality in Appendix B.2. □

Remark 3.2 (Extension to arbitrary recomposition functions)**.** In the instance giving Theorem 3.1, either utility matrix gave all voters identical stakes. If $s(\mathbf{u}_i)$ is equal across voters, then for any recomposition function r, $r(s(\mathbf{u}_i))$ will be equal across voters. *Observation 2* still holds, and the lower bound applies. Given that we will find this lower bound to be tight, it implies that in the worst case, applying a recomposition function other than the identity function will not improve the s-distortion of any deterministic voting rule, for any stakes function s.

Remark 3.3 (Connection to existing results)**.** Theorem 7 of Amanatidis et al. [2021] shows that any single value query of utilities (where the voting mechanism can directly ask agents about any one of their utilities) can enable at best $\Omega(m)$ distortion. Our lower bound in Theorem 3.1 generalizes this lower bound, showing that *any system of queries yielding the value of a scalar-valued stakes function*, when paired with a deterministic voting rule, can achieve at best $\Omega(m)$ distortion.[4]

3.1.2 Upper bound for all f, s. Now, we will prove upper bound on the s-distortion of *any* deterministic voting rule f, for *any* stakes function s. To reason about all voting rules f and stakes functions s at once, we must determine: *Given any pair of s, f, what properties of s and f will lead to low distortion?* We now introduce two such properties. First, β_f is the minimum fraction of voters that must rank the winner by f in first position:

$$\beta_f := \min_{\mathbf{h}\in\Delta(S_m)} \textstyle\sum_{\pi\in S_m} h_\pi \cdot \mathbb{I}\{\pi(1) = f(\mathbf{h})\}.$$

Second, κ-upper(s) and κ-lower(s) measure the extent to which s can over- or under-estimate $\max(\mathbf{u})$, respectively:

$$\kappa\text{-upper} := \sup_{\mathbf{u}} \frac{s(\mathbf{u})}{\max(\mathbf{u})}, \quad \kappa\text{-lower}(s) := \inf_{\mathbf{u}} \frac{s(\mathbf{u})}{\max(\mathbf{u})}.$$

While bounds in terms of other properties of s, f are conceivable, these quantities will permit optimal upper bounds.

In terms of these quantities, Theorem 3.4 gives an upper bound on the s-distortion for any s and any deterministic f. The proof relies on the insight that β_f and the κ values are linked: β_f lower bounds how often the winner is ranked first, while the κ's links the stakes and maximum utility of any voter who ranks the winner first. This connection implies a lower-bound on the social welfare of the winner.

Theorem 3.4 (**upper bound**). *For all s and deterministic f,*

$$\mathsf{dist}^{\mathbf{s}}(f) \le \beta_f^{-1} \cdot \kappa\text{-upper}(s) \,/\, \kappa\text{-lower}(s).$$

Proof. Fix an instance U, a stakes function s, and a deterministic rule f. Let $a' = f(\mathsf{hist}^s(U))$ be the winner of the s-proportional election. First, we have that the social welfare of any alternative a is upper-bounded:

$$\mathsf{sw}(a, U) \le \sum_{i\in[n]} \max(\mathbf{u}_i) \le \sum_{i\in[n]} s(\mathbf{u}_i) \,/\, \kappa\text{-lower}(s). \tag{1}$$

Now, let $N_{a'}$ be the set of voters who rank a' first. All $i \in N_{a'}$ must have at least some utility for a':

$$u_i(a') = \max_a u_i(a) \ge s(\mathbf{u}_i)/\kappa\text{-upper}(s). \tag{2}$$

Also, since a' is the winner, $N_{a'}$ composes at least a β_f fraction of the stakes-proportional electorate:

$$\textstyle\sum_{i\in N_{a'}} s(\mathbf{u}_i) \,/\, \sum_{i\in[n]} s(\mathbf{u}_i) \ge \beta_f.$$

[4]Note: a necessary step in showing that our lower bound subsumes theirs is arguing that our lower bound actually applies to *any stakes-dependent electoral recomposition*, not just proportional recomposition, which we do.

This fact, combined with Equation (2), gives that

$$\mathrm{sw}(a', U) \geq \sum_{i \in N_{a'}} u_i(a') \geq \sum_{i \in N_{a'}} s(\mathbf{u}_i) \,/\, \kappa\text{-upper}(s) \geq \beta_f \sum_{i \in [n]} s(\mathbf{u}_i) \,/\, \kappa\text{-upper}(s).$$

Combining this with Equation (1) and denoting the maximum welfare alternative by a^*, we obtain that

$$\mathrm{dist}_U^s(f) = \frac{\mathrm{sw}(a^*, U)}{\mathrm{sw}(a', U)} \leq \beta_f^{-1} \cdot \frac{\kappa\text{-upper}(s)}{\kappa\text{-lower}(s)}. \qquad \square$$

Remark 3.5. Theorem 3.4 also holds if in the definitions of κ-upper(s) and κ-lower(s), max is replaced with range. This is because the worst-case distortion can always be realized by instances where every voter has minimum utility 0, in which case max = range. We prove this in Appendix B.3.

3.1.3 Optimality of Plurality and s =max. In our first extension of these results, we identify a voting rule-stakes function pair that attains the best possible s-distortion m, matching our lower bound in Theorem 3.1. We do this by minimizing the upper bound in Theorem 3.4, which amounts to choosing s and f to minimize both κ-upper$(s)/\kappa$-lower(s) and β_f^{-1}.

It is easy to see that $s = \max$ achieves the minimal value κ-upper$(s)/\kappa$-lower$(s) = 1$. For β_f, the answer is more subtle; in Appendix B.4, we prove the following lemma, which shows that the maximal attainable value is $\beta_f = 1/m$, and that this maximum is achieved by Plurality.

Lemma 3.6. *For any deterministic voting rule f, $\beta_f \leq 1/m$, and $\beta_{\text{Plurality}} = 1/m$.*

With this lemma in hand, we now apply Theorem 3.4 to conclude the following upper bound, which matches the lower bound in Theorem 3.1.

Proposition 3.7. *$\mathrm{dist}^{\max}(\text{Plurality}) \leq m$.*

The above result suggests using Plurality with $s = \max$ as a promising choice. However, in some motivating contexts — e.g., where stakes-proportionality arises from voters' behavior — we may not be able to control which stakes function is used. To characterize the s-distortion of Plurality, we show that it is essentially tight with respect to Theorem 3.4 for all s, except that κ-lower is replaced with $\tilde{\kappa}$-lower, in which the supremum is taken over only utility vectors $\mathbf{u}$ with the same first and second entry. See Appendix B.5 for the formal definition of $\tilde{\kappa}$-lower and subsequent proof.

Proposition 3.8. *For all s, $\mathrm{dist}^s(\text{Plurality}) \geq (m-1) \cdot \frac{\kappa\text{-upper}(s)}{\tilde{\kappa}\text{-lower}(s)}$.*

Remark 3.9 (Connection to existing results)**.** Theorem 1 of Amanatidis et al. [2021] shows that their voting mechanism 1-PRV — which is equivalent to Plurality under stakes-proportionality with respect to max — gives distortion $O(m)$. This result corresponds to our upper bound on the max-distortion of Plurality, proven via Proposition 3.7.

3.1.4 Beyond Plurality. To understand the s-distortion of other deterministic voting rules f, we first observe that when $\beta_f = 0$, there is an unbounded gap between the lower bound in Theorem 3.1 and the upper bound in Theorem 3.4. Unfortunately, the s-distortion for such f is indeed unbounded — a finding that is practically significant because, as we prove in Appendix B.6, most popular voting rules have $\beta_f = 0$.

Proposition 3.10. *For all stakes functions s and all deterministic rules f with $\beta_f = 0$, $\mathrm{dist}^s(f) = \infty$.*

Proof. Let f satisfy $\beta_f = 0$, and fix a histogram $\mathbf{h}$ in which the winner $f(\mathbf{h})$ is never ranked first. Then, set the underlying U to realize this histogram while setting each voter's ordered utility vector to $\mathbf{1}_1\mathbf{0}_{m-1}$. Since the winner is never ranked first, it must get 0 average utility. Since each voter gives their respective first-ranked alternative utility 1, at least one alternative must have at least $1/m$

average utility; thus, $\mathsf{dist}_U(f) = \infty$ is unbounded. Because all voters have identical utility vectors, for all s, $\mathsf{hist}(U) = \mathsf{hist}^s(U) \implies f(\mathsf{hist}(U)) = f(\mathsf{hist}^s(U)) \implies \mathsf{dist}_U(f) = \mathsf{dist}^s_U(f) = \infty$. □

Proposition 3.7 and Proposition 3.10 together point to Plurality-like rules as uniquely promising when stakes are accounted for. This may seem strange, as Plurality is often considered a "bad rule" due to its lack of expressiveness. One positive interpretation of this finding is that Plurality, when stakes are accounted for, actually accounts for the most critical information; this is good news, as Plurality-like voting methods are widely used. Another possibility is that the ranking-based ballot format is insufficiently expressive; as we discuss in Section 6, our model and approach extend easily to richer ballot formats.

It is also important to acknowledge that although $\beta_f = 0$ for many non-Plurality voting rules, this is a *worst case* result. It could very well be that in typical instances, the winners chosen by other voting rules are often ranked first by many voters. In this case, our impossibility in Proposition 3.10 would be quite pessimistic. This possibility motivates beyond-worst-case analysis, and/or simple sufficient conditions under which a broader range of voting rules can make use of stakes information. We leave these directions to future work.

3.2 Randomized voting rules

Next, we give an analogous but brief analysis of randomized rules. First, we lower-bound the s-distortion across all f and all 1-homogeneous s, characterizing what s-distortion stakes-proportionality will permit. For comparison, randomized rules have at best $\Omega(m)$ distortion without accounting for stakes (see Appendix B.10 for details).

Theorem 3.11 (**lower bound**). *For all 1-homogeneous s, randomized f,* $\mathsf{dist}^s(f) \geq \frac{\sqrt{m}}{10+3\log m}$.

Proof sketch. The construction of this lower bound is rather intricate, and its full proof is deferred to Appendix B.7. The main idea is to identify indices of the utility vector over which large gaps in utilities have the smallest effects on the stakes. More formally, we try to identify a small $z \in [m]$ such that the following quantity is upper-bounded:

$$\frac{s(\mathbf{1}_{z+1}\mathbf{0}_{m-(z+1)})}{s(\mathbf{1}_z\mathbf{0}_{m-z})}.$$

Then, we can exploit this lack of sensitivity of the stakes function to drive large gaps in alternatives' utilities over this index. If such a small z doesn't exist, we know that placing utility gaps over all early indices of a utility vector will create a lower-bounded gap in stakes, which we can then exploit via a different construction. As in previous lower bounds, in the instance we design, all voters have identical stakes. □

Our next result proves that a known randomized rule, Stable Lottery, paired with one of a few stakes functions s, achieves $O(\sqrt{m})$ s-distortion, matching our lower bound up to a log factor. This rule was originally introduced with the goal of achieving low distortion in the distinct setting assuming normalized utilities; there, Stable Lottery was shown to achieve distortion $O(\sqrt{m})$ when utilities were restricted to have have unit sum, i.e., $\sum_{a\in[m]} u_i(a) = 1$ for all i [Ebadian et al., 2022]. The theorem is proven via Lemma 3.13, which shows a connection between our model and the popular normalized utilities model that may be of independent interest.

Theorem 3.12 (**upper bound**). *For $s \in \{sum, max, range\}$,* $\mathsf{dist}^s(\text{Stable Lottery}) \in O(\sqrt{m})$.

Proof. We prove this upper bound by connecting stakes information to the popular normalized utility model, in which it is assumed that voters' utilities are normalized to sum to 1 [Caragiannis et al., 2017, Caragiannis and Procaccia, 2011, Procaccia and Rosenschein, 2006] (or, in recent work, have *maximum* utility 1 [Ebadian et al., 2022]). Placed within our model, observe that such assumptions amount to assuming that voters have *identical stakes* as measured by sum (resp. max).[5] Of course, one could assume this normalization with respect to *any* stakes function s; we call this general class of assumptions *s-unit-stakes* assumptions.

We now show a surprising equivalence: Assuming sum-unit stakes is *equivalent*, from a distortion perspective, to the sum-distortion (i.e., the distortion achievable under s-proportionality when $s = \text{sum}$). In fact, we show this correspondence to hold for *any 1-homogeneous stakes function*, as stated informally in Lemma 3.13.

Lemma 3.13 (informal). *For all 1-homogeneous s, the distortion of f under the s-unit-stakes assumption is equivalent to its s-distortion.*

The bidirectional reduction that proves this lemma is pictured in Figure 1. We prove it formally over two theorems in Appendix B.8: Theorem B.5 handles rational-valued histograms, and Theorem B.7 extends the claim to real-valued histograms (at the cost of a mild technical condition on the voting rule). With this reduction in hand, we can now transfer lower and upper bounds

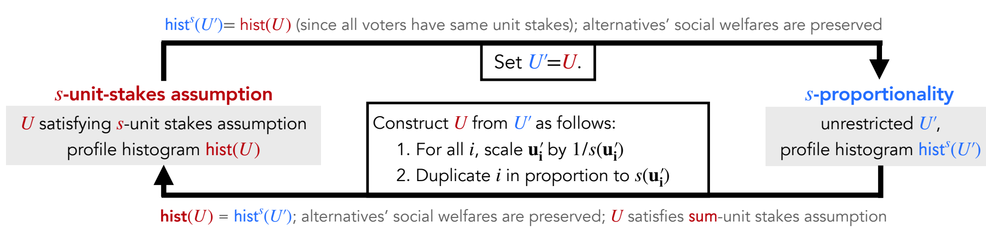

Fig. 1. Constructions for reducing between the s-unit stakes assumption (existing model) and s-proportionality (our model).

between the s-unit stakes setting and ours. We conclude the proof of Theorem 3.12 by transferring bounds on Stable Lottery proven under the sum-, max-unit stakes assumptions [Ebadian et al., 2022, Theorem 3.4]. The case of range requires a minor technical extension — see Appendix B.9 for details. □

Implications for the s-unit stakes literature. This reduction proves a formal connection between two distinct branches of the literature: distortion under normalized utilities, and distortion with auxiliary utility information. Using the fact that this reduction permits transferring bounds across these settings, in Appendix A.1 we illustrate how our results recover — and even sometimes strengthen — existing results, often via simpler arguments.

Beyond existing results, we can also use our reduction to immediately derive new results for the s-unit stakes assumption setting. Although there is a vast space of possible *s-unit-stakes* assumptions (one per possible stakes function s), the literature has explored few; to our knowledge, only max and sum. This inspires the open question: *would assuming unit stakes with respect to a different stakes function permit better distortion bounds?* Our results immediately close this question for deterministic rules:

Corollary 3.14 (of Lemma 3.13). *For all 1-homogeneous stakes functions s and all deterministic voting rules f, the best achievable distortion of any f under any s-unit stakes assumption is $m - 1$*

[5]This interpretation is already informative, as it raises substantive questions about whether these widely-used restrictions on the utilities are likely to hold in practice, where people may have dramatically differing stakes.

(Theorem 3.1), and this is achieved by assuming max-unit stakes or range-unit stakes and using the rule Plurality (Proposition 3.7, Remark 3.5).

Our results also close this question within a factor of $\log(m)$ for randomized rules:

Corollary 3.15 (of Lemma 3.13)**.** *For all 1-homogeneous stakes functions s and all randomized voting rules f, the best achievable distortion of any f under any s-unit stakes assumption is at least $\Omega(\sqrt{m}/\log(m))$ (Theorem 3.11), and this distortion is achieved within a $\log(m)$ factor by assuming* max-unit stakes, range-unit stakes, or sum-unit stakes and using the rule Stable Lottery (Theorem 3.12).

4 WHAT IF WE HAVE *APPROXIMATE* STAKES INFORMATION?

In many cases, we may know voters' stakes only coarsely, at the level of orders of magnitude. To understand what is possible in this case, we now extend our results from previous sections to the case where our estimates of voters' stakes are incorrect.

Let us formally define these errors. Suppose we achieve s-proportionality according to an incorrect estimate of each voter i's stakes $\hat{s}(\mathbf{u}_i) := \delta_i\, s(\mathbf{u}_i)$, where $\delta_i \geq 1$ is the factor by which we overestimate i's stakes.[6] Let $\boldsymbol{\delta} := (\delta_i | i \in [n])$ be the vector of all such errors. Given U and $\boldsymbol{\delta}$, we denote the $\boldsymbol{\delta}$-approximately stakes-proportional histogram as $\mathsf{hist}^{\boldsymbol{\delta},s}_\pi(U)$, with π-th entry

$$\mathsf{hist}^{\boldsymbol{\delta},s}_\pi(U) := \frac{\sum_{i\in[n]} \hat{s}(\mathbf{u}_i) \cdot \mathbb{I}(\pi_i = \pi)}{\sum_{i\in[n]} \hat{s}(\mathbf{u}_i)}.$$

For $\delta := \max_i \delta_i$, the δ, s-distortion of f is then given as

$$\mathsf{dist}^{\delta,s}(f) = \sup_{n\geq 1,\ U,\ \boldsymbol{\delta}\in[1,\delta]^n} \frac{\max_a \mathsf{sw}(a,U)}{\mathbb{E}[\mathsf{sw}(f(\mathsf{hist}^{\boldsymbol{\delta},s}(U),U)]}.$$

Note that for fixed δ, by this definition, $\boldsymbol{\delta} \in [1,\delta]^n$ is chosen adversarially. We now prove strong, instance-wise robustness to such errors.

Theorem 4.1. *For all f, 1-homogeneous s, U, and $\delta \geq 1$, $\mathsf{dist}^{\delta,s}_U(f) \leq \delta\, \mathsf{dist}^s_U(f)$.*

The intuition is simple: since s is 1-homogeneous, mis-estimating i's stakes by up to δ is the same as overestimating voters' *utilities* by up to δ. Such overestimates can change the distortion by at most a δ factor. We now formalize this intuition.

Proof of Theorem 4.1. Fix a utility matrix U, a 1-homogeneous stakes function s, and an error vector $\boldsymbol{\delta} \in [1,\delta]^n$. Let $\tilde{U}$ be the utility matrix where voter i's utility vector is scaled by a factor of δ_i, i.e., $\tilde{\mathbf{u}}_i = \delta_i \mathbf{u}_i$. Then, since s is 1-homogeneous, we have that $s(\tilde{\mathbf{u}}_i) = \delta_i s(\mathbf{u}_i)$, and therefore $\mathsf{hist}^s(\tilde{U}) \equiv \mathsf{hist}^{\boldsymbol{\delta},s}(U)$, directly implying that $f(\mathsf{hist}^s(\tilde{U})) \equiv f(\mathsf{hist}^{\boldsymbol{\delta},s}(U))$. Moreover, for every alternative a, it holds that $\mathsf{sw}(a,\tilde{U}) \in [\mathsf{sw}(a,U), \delta\mathsf{sw}(a,U)]$. It follows that

$$\mathbb{E}[\mathsf{sw}(f(\mathsf{hist}^s(\tilde{U})),\tilde{U})] \leq \delta \cdot \mathbb{E}[\mathsf{sw}(f(\mathsf{hist}^{\boldsymbol{\delta},s}(U)),U)],$$

from which we deduce that

$$\frac{\max_a \mathsf{sw}(a,U)}{\mathbb{E}[\mathsf{sw}(f(\mathsf{hist}^{\boldsymbol{\delta},s}(U)),U)]} \leq \delta \frac{\max_a \mathsf{sw}(a,U)}{\mathbb{E}[\mathsf{sw}(f(\mathsf{hist}^s(\tilde{U})),\tilde{U})]} = \frac{\max_a \mathsf{sw}(a,U)}{\max_a \mathsf{sw}(a,\tilde{U})} \cdot \delta \cdot \frac{\max_a \mathsf{sw}(a,\tilde{U})}{\mathbb{E}[\mathsf{sw}(f(\mathsf{hist}^s(\tilde{U}),\tilde{U})]}$$
$$\leq \delta \cdot \mathsf{dist}^s(f).$$

Taking suprema on the left hand side then completes the proof. □

[6] We can realize any type of errors with $\delta \geq 1$, because the weighting of the resulting electorate is relative.

5 WHAT IF WE HAVE *NO* STAKES INFORMATION?

In this section, we are interesed in the case where voters may have extremely different stakes (in which case our results dictate that accounting for them is important), but either (1) we cannot even coarsely guess them, or (2) we cannot explicitly re-weight votes without public backlash. In this section, we propose a mechanism design concept for *revealing voters' stakes through their behavior.* We propose a class of mechanisms, called *multi-issue mechanisms*, based on a simple idea: if voters must decide how to allocate voting power over *multiple entire elections*, they will exert influence in the decisions that most affect them, thereby revealing their stakes. We define this class of mechanisms here, where *(MD)* indicates a decision by the mechanism designer:

Multi-issue mechanisms:

Setup. Each voter $i \in [n]$ is presented with a slate of k entire elections *(MD)*, where each election ℓ is over its own set of alternatives A^ℓ. Voter i has utility vector $\mathbf{u}_i^\ell$ in election ℓ.

Voting. To vote, each voter i submits to each election two things: a ballot of some format *(MD)*, and a scalar *weight* $w_i^\ell \in [0, 1]$ describing the weight i wishes to place on election ℓ. These weights are restricted such that $\sum_\ell c(w_i^\ell) \leq 1$ for all i, where $c : \mathbb{R} \to \mathbb{R}^+$ is a cost function *(MD)* describing how much a voter is charged per unit of weight placed on an election.

Aggregation. Each voter i's ballot in each election ℓ is weighted by w_i^ℓ, and a voting rule *(MD)* is then used to aggregate these weighted ballots into a winner.

We next perform a proof-of-concept analysis of a multi-issue mechanism. We emphasize that this analysis is meant to build intuition rather than to be the final word on the design of multi-issue mechanisms. In particular, the analysis will illustrate (1) how voters' individual incentives can reveal their stakes according to a natural stakes function, (2) what assumptions are required for analyzing such a mechanism, and (3) how our results from Section 3 and Section 4 can be applied to bound the distortion of such a mechanism. Then, with the intuition from this analysis in hand, in Section 6.1, we will discuss generalizations of this mechanism and connections to the existing mechanisms *storable votes* [Casella, 2005] and *quadratic voting* [Posner and Weyl, 2015], which use similar types of trade-offs to elicit richer utility information.

To fully define our multi-issue mechanism, we must instantiate the four aspects left up to the mechanism designer (those flagged with *(MD)*): the elections to be included, the ballot format, the cost function, and and the voting rule.

Elections to be included. We allow the inclusion of an arbitrary number of elections $k \geq 2$. For ease of exposition, we assume each election ℓ contains the same number of alternatives m (this can easily be achieved by simply adding dummy alternatives).

Ballot format. Voters "vote" by submitting complete rankings as was assumed throughout the paper, in order to apply our results to analyze the mechanism.

Cost function. We let $c : \mathbb{R}^k \mapsto \mathbb{R}$ be quadratic,[7] so that given $\mathbf{x} := (x_i | i \in [k]), c(\mathbf{x}) = \sum_{i \in [k]} x_i^2$.

Voting rule. We will aggregate votes via Plurality, guided by our results showing its optimality among deterministic rules.

For the purposes of analyzing our mechanism, we must also define with what motives — and based on what information — voters translate their utilities into actions (votes). Our assumptions prioritize simplicity.

[7]One may wonder if there is a correspondence between this mechanism and Quadratic Voting (QV). Technically, the cost quadraticity serves the same purpose, but our mechanism is over *multiple elections*, while QV occurs within a single election.

Assumption 5.1. *We assume that voters satisfy the following properties.*

- *Honest:* voters will submit rankings that are true to their underlying utilities.
- *Oblivious*: voters have uniform priors over all sets of $n-1$ other rankings in every election, i.e., they make the *impartial culture assumption* [Gehrlein, 2002].
- *Not mathematicians:* voters believe their probability of pivotality in election ℓ increases linearly in w_i^ℓ, the weight they place on their vote in that election.[8]
- *Utility-maximizing:* voters submit weights across elections to maximize their total expected individual utility (where the expectation is over the randomness of others' votes).

Finally, we need to introduce formal notation that will allow us to analyze multi-issue mechanisms. In each election $\ell \in [k]$, the underlying $n \times m$ utility matrix is denoted U^ℓ; correspondingly, voters' utility vectors are $\mathbf{u}_i^\ell$ and individual utilities are $u_i^\ell(a)$. For a given stakes function s, we use shorthand $\boldsymbol{s}_i := s(\mathbf{u}_i^\ell)_{\ell \in [k]}$ to summarize i's stakes across all k elections.

Because we use Plurality, we are concerned only with each voter's first-ranked alternative, so we will consider histograms that are indexed by $a \in m$ rather than $\pi \in S_m$. Let $b_i^\ell \in [m]$ be i's first-ranked alternative in election ℓ; abusing notation, we will represent i's vote in this election using $\boldsymbol{e}_{b_i^\ell}$, the m-length alternative-indexed basis vector with a 1 at the b_i^ℓ-th index. We can then represent i's weighted vote in election ℓ as the weighted basis vector $w_i^\ell \boldsymbol{e}_{b_i^\ell}$. Let $\boldsymbol{w}_i := (w_i^\ell | \ell \in [k])$ be the weights i submits across elections. We let the resulting histogram in election ℓ be $\mathbf{h}^\ell := \sum_{i \in [n]} w_i^\ell \boldsymbol{e}_{b_i^\ell} / \sum_{i \in [n]} w_i^\ell$. Let $\mathbf{h}_{-i}^\ell$ be the histogram not including i's vote, i.e., $\mathbf{h}_{-i}^\ell := \sum_{i \in [n] \setminus \{i\}} w_i^\ell \boldsymbol{e}_{b_i^\ell} / \sum_{i \in [n]} w_i^\ell$. We write $\mathbf{h} \sim \mathcal{I}$ to denote a histogram that reflects a profile drawn from the Impartial Culture model.

Now, we apply results from Sections 3 and 4 to show that although any election encompassed by the mechanism, if run in isolation, could have had unbounded distortion, within this mechanism all elections are guaranteed to have distortion at most order m^2 (Theorem 5.4). This analysis proceeds in 3 steps. In **Step 1**, we show how voters' stakes are revealed in how they allocate weights across elections, and we characterize the particular stakes function that arises from their behavior. Then, in **Step 2**, we show that this stakes function is quite natural, and per our previous results, behaves similarly enough to range to permit near optimal distortion when paired with Plurality. Finally, in **Step 3**, we characterize the distortion of each election in the mechanism, which reveals an important feature of our choice of the k elections around which we design the mechanism.

Step 1: A natural stakes function s^* arises from voter behavior. First, we show that the stakes function arising from voters' utility-maximizing behavior is simple and intuitive:

$$s^*(\mathbf{u}) := \mathsf{max}(\mathbf{u}) - \frac{\mathsf{sum}(\mathbf{u}) - \mathsf{max}(\mathbf{u})}{m-1},$$

the gap between their maximum utility and their average utility for the other alternatives.

Lemma 5.2. *In each election ℓ, each voter i weights their vote by $\hat{w}_i^\ell = s^*(\mathbf{u}_i^\ell) \,/\, \|\boldsymbol{s}_i^*\|_2$.*

Proof. Fix an i, whose behavior in the mechanism we will analyze. Let the alternatives in each election ℓ be A^ℓ, and for all $\ell \in [k]$, let b_i^ℓ be i's favorite alternative in the ℓ-th election. Now, define the following events, some which depend on $\boldsymbol{w}_i$:

[8]While this is technically inconsistent with their uniform priors over others' votes, we see it as at least as reasonable as assuming voters compute their precise probabilities of pivotality: even under uniform priors, computing the exact marginal increase in probability of pivotality per unit of weight is a combinatorial calculation that we cannot expect voters to do. In light of this, linearity is a natural assumption.

- X^ℓ: b_i^ℓ wins among only the votes in $\mathbf{h}_{-i}^\ell$ (and thus also wins with i's ranking weighted by w_i^ℓ added to the profile).
- $Z^\ell(\boldsymbol{w}_i)$: Some $a \neq b_i^\ell$ wins among only the votes in $\mathbf{h}_{-i}^\ell$ by a margin small enough that i's vote for b_i^ℓ is *pivotal.*
- $\neg X^\ell \wedge \neg Z^\ell(\boldsymbol{w}_i)$: Some alternative $a \neq b^\ell$ wins among the votes in $\mathbf{h}_{-i}^\ell + w_i^\ell \boldsymbol{e}_{b_i^\ell}$ (i.e., regardless of whether i votes for b_i^ℓ with weight w_i^ℓ).

Note that these events are mutually exclusive. Now, we compute voter i's expected reward in terms of these quantities:

$$\mathbb{E}_{\mathbf{h}_{-i}\sim\mathcal{I}}\left[\sum_{\ell\in[k]} u_i(\textsc{Plurality}(\mathbf{h}_{-i} + w_i^\ell \boldsymbol{e}_{b_i^\ell}))\right]$$
$$= \sum_{\ell\in[k]}\left(\left(\Pr[X^\ell] + \Pr\left[Z^\ell(\boldsymbol{w}_i)\right]\right)\cdot u_i(b_i^\ell) + Pr\left[\neg X^\ell \wedge \neg Z^\ell(\boldsymbol{w}_i)\right]\cdot \sum_{a\in A^\ell\setminus\{b_i^\ell\}} \frac{u_i(a)}{m-1}\right)$$

Unpacking this expression, if either event X^ℓ or Z^ℓ occurs, i gets the utility associated with b_i^ℓ. If neither occurs, i expects some other alternative to win, and has uniform priors over other voters' rankings; thus their expected utility in this event is the average of all alternatives in $A^\ell \setminus \{b_i^\ell\}$. Again using i's uniform priors over other voters' behavior, $Pr[X^\ell] = 1/m$. Simplifying,

$$= \sum_{\ell\in[k]}\left(\left(1/m + \Pr[Z^\ell(\boldsymbol{w}_i)]\right) u_i(b_i^\ell) + \left(1 - 1/m - Pr[Z^\ell(\boldsymbol{w}_i)]\right)\cdot \sum_{a\in A^\ell\setminus\{b_i^\ell\}} \frac{u_i(a)}{m-1}\right)$$
$$= \sum_{\ell\in[k]}\left(\sum_{a\in A^\ell}\frac{u_i(a)}{m} + Pr[Z^\ell(\boldsymbol{w}_i)]\left(u_i(b_i^\ell) - \sum_{a\in A^\ell\setminus\{b_i^\ell\}}\frac{u_i(a)}{m-1}\right)\right).$$

The first term of the summand is fixed in the instance, so it is not decision relevant. Then, we have deduced that

$$\max_{\boldsymbol{w}_i}\mathbb{E}_{\mathbf{h}_{-i}\sim\mathcal{I}}\left[\sum_{\ell\in[k]} u_i(\textsc{Plurality}(\mathbf{h}_{-i} + w_i^\ell \boldsymbol{e}_{b_i^\ell}))\right]$$
$$= \max_{\boldsymbol{w}_i}\sum_{\ell\in[k]} Pr[Z^\ell(\boldsymbol{w}_i)]\left(u_i(b_i^\ell) - \sum_{a\in A^\ell\setminus\{b_i^\ell\}}\frac{u_i(a)}{m-1}\right)$$
$$= \max_{\boldsymbol{w}_i}\sum_{\ell\in[k]} Pr[Z^\ell(\boldsymbol{w}_i)]\left(\mathsf{max}(\mathbf{u}_i^\ell) - \frac{\mathsf{sum}(\mathbf{u}_i^\ell) - \mathsf{max}(\mathbf{u}_i^\ell)}{m-1}\right)$$
$$= \max_{\boldsymbol{w}_i}\sum_{\ell\in[k]} Pr[Z^\ell(\boldsymbol{w}_i)]\cdot s^*(\mathbf{u}_i^\ell)$$

Finally, by Assumption 5.1, voters act as though their probability of pivotality increases linearly with each additional unit of weight placed behind their vote. Applying this assumption,

$$= \max_{\boldsymbol{w}_i}\sum_{\ell\in[k]} w_i^\ell\cdot s^*(\mathbf{u}_i^\ell).$$

Then, subject to to the constraint that each voter has total weight 1, the final problem voters are solving is as follows, where $\hat{\boldsymbol{w}}_i$ describes their optimal weighting across elections:

$$\hat{\boldsymbol{w}}_i := \arg\max_{\boldsymbol{w}_i} \sum_{\ell \in [k]} w_i^\ell \cdot s^*(\mathbf{u}_i^\ell) \quad \text{s.t.} \quad \sum_{\ell \in [k]} (w_i^\ell)^2 \leq 1. \tag{3}$$

We can compute the optimizer of this program by simply projecting the vector $\boldsymbol{s}_i^*$ onto the unit sphere, concluding that

$$\hat{\boldsymbol{w}}_i = \frac{\boldsymbol{s}_i^*}{\|\boldsymbol{s}_i^*\|_2}. \qquad \square$$

Step 2: This stakes function s^* is nearly optimal. Next, using our previous distortion bounds, we find that s^* is *nearly optimal* across all stakes functions when paired with Plurality:

Lemma 5.3. $\mathsf{dist}^{s^*}(\text{Plurality}) \leq m^2$.

Proof. We will use the alternative definitions of κ-upper and κ-lower from Remark 3.5, where max is replaced with range. Characterizing these κ values, $\kappa\text{-upper}(s^*) = 1$, realized by the vector $\mathbf{u} = \mathbf{1}_1\mathbf{0}_{m-1}$, and $\kappa\text{-lower}(s^*) = 1/(m-1)$, realized by the vector $\mathbf{u} = \mathbf{1}_{m-1}\mathbf{0}_1$. Using that $\beta_{\text{Plurality}} = 1/m$ (Lemma 3.6), Theorem 3.4 gives us that $\mathsf{dist}^{s^*}(\text{Plurality}) \leq m \cdot \frac{1}{1/(m-1)} \leq m^2$. $\square$

Step 3: Bounding the distortion of the mechanism. Finally, we bound the distortion in each election within the mechanism. The remaining issue to deal with is that although each *individual voter* will spread their votes across elections in proportion to their stakes (Lemma 5.2), if some voters have substantially higher *total stakes across the k elections*, the uniform budgets across voters will under-count these voters' stakes. It may seem that we are back where we started — where uniform voting power fundamentally cannot account for stakes. However, unlike before, we have another lever at our disposal: *the design of our slate of k elections*. Specifically, we can choose this slate of elections such that all voters are affected to a relatively high degree by at least one election. While we cannot hope for this approach to perfectly equalize voters' total stakes, it may bring them *closer*, e.g., within a factor of δ. Then, per Theorem 4.1, all elections in our mechanism will have distortion at most δm^2.

Theorem 5.4. *Fix a $\delta \geq 1$ such that for all pairs of voters $i, i' \in [n]$, $\|\boldsymbol{s}_i^*\|_2 / \|\boldsymbol{s}_{i'}^*\|_2 \leq \delta$. Then, for all $\ell \in [k]$,*

$$\mathsf{dist}_{U^\ell}^{s^*}(\textit{Plurality}) \leq \delta m^2.$$

Proof. Let $\alpha := \max_{i \in [n]} \|\boldsymbol{s}_i^*\|_2$ be the maximum total stakes of any voter. Fix an $\ell \in [k]$. Then, per Lemma 5.2, for all i, and for some $\delta_i \in [1, \delta]$, we have that $w_i^{*\ell} = s^*(\mathbf{u}_i^\ell) / \|\boldsymbol{s}_i^*\|_2 = \delta_i \cdot s^*(\mathbf{u}_i^\ell)/\alpha$. Since α is constant across voters, we get s^*-proportionality in election ℓ with respect to the mis-estimate of i's stakes $\tilde{s}^*(\mathbf{u}^\ell) = \delta_i \cdot s^*(\mathbf{u}^\ell)$. This is the precondition of Theorem 4.1; using this and Lemma 5.3,

$$\mathsf{dist}_{U^\ell}^{\delta, s^*}(\text{Plurality}) \leq \delta\, \mathsf{dist}_{U^\ell}^{s^*}(\text{Plurality}) \leq \delta m^2. \ \square$$

6 DISCUSSION

We conclude with an in-depth discussion of questions raised by multi-issue mechanisms, followed by opportunities for accounting for stakes in emerging democratic paradigms.

6.1 The multi-issue mechanism design space & connections to other mechanisms

While we focused on ranking-based voting, the same mechanistic approach we introduced for multi-issue mechanisms could be used with essentially any election format: simply place multiple elections on the same ballot, across which voters must trade off a total allotment of voting power. In fact, there is a known mechanism that roughly does this: *storable votes* [Casella, 2005], which allows voters to save up votes over a sequence of elections and spend them on the later elections they care about most. Like the mechanism above, storable votes is just one within a massive design space, whose many levers we explore below. As we go, we point out a wealth of open questions.

Slate of issues. As captured by δ in our analysis above, the performance of a multi-issue mechanism relies on choosing a slate of issues that *roughly balance* voters' total stakes. In some cases this will not be hard: consider a case where we want to decide a design aspect of our public transit systems, and want to ensure we sufficiently account for the interests of those who cannot afford cars. We might choose our second issue to be where to fix potholes — a decision that will *primarily* affect those who drive. In this case, voters who do/do not have cars will have high stakes in opposite elections, getting closer to balanced total stakes than in either election alone.

However, this gets complicated quickly when we want to consider more than two groups of stakeholders. For example, suppose there is a third group in our example above — people who drive, but also live by candidate transit stops. These people may have high stakes in *both decisions*, due to their potential to avoid pothole damage to their cars *and* their potential to be affected by loud public transportation outside their house. To balance stakes across these voters too, we need a third issue, which may in turn imbalance the stakes of the first two groups.

To state this question of issue design formally: Suppose there is a universe of potential elections L and a set of voters $[n]$, where for each election $\ell \in L$, each voter i has stakes $s(\mathbf{u}_i^\ell)$. For any given $K \subseteq L$, let $\mathbf{s}_i^K := (s(\mathbf{u}_i^\ell)|\ell \in K)$. Then, for any given $\delta \geq 1$, we want to know: *what is the smallest slate of elections K that ensures that for all $i, i' \in [n]$, $\|\mathbf{s}_i^*\|_2 \,/\, \|\mathbf{s}_{i'}^*\|_2 \leq \delta$?* As illustrated by our analysis above, the precise stakes function and method of totaling stakes for which we want to achieve this bound is mechanism-dependent.

Ballot format. In our paper, we assumed voters submit their preferences as complete rankings. However, there are many other ballot formats that could reveal richer information about voters' utilities. One promising candidate is the ballots used in Quadratic Voting (QV) [Posner and Weyl, 2015], where a voter has a budget of total voting power to allocate *over the alternatives in a single election*, and they are charged quadratically per unit of weight they place on a given alternative. For example, in an election over alternatives a, b, c where each voter has 10 votes total, they could place 3 on alternative a, 1 on b, and 0 on c to spend their total budget of $3^2 + 1^2 + 0^2 = 10$.

Notice that this is different from how we use quadraticity in our mechanism: we charge voters quadratically for power allocated toward a given *election*, while QV considers a single election and charges voters for power allocated toward a single *alternative*. However, the quadraticity works for the same reason, encouraging voters to spread their power over elections or alternatives *in proportion* to their stakes or relative utilities, respectively. In summary, QV ballots — at least under the (strong) assumptions made by Posner and Weyl [2015] — allow the recovery of information about voters' *relative utilities* (i.e., i likes a 3 times as much as b), while rankings lose this information, thereby losing m distortion even when stakes are known. To design a multi-issue mechansim using QV ballots, then, one must answer: *how can we set up voting budgets to recover stakes-proportionality when voters are charged quadratically per unit of votes for a single alternative* and *weight toward a single election?*

Designing new ballots will open up the possible space of voting rules — an enticing prospect, given that with ranking-based ballots, our results show that the space of rules where stakes information is helpful is very limited.

Voter model. Ultimately, the success of a multi-issue mechanism depends on how voters will *behave* within it. This depends on several things, including (1) voters' beliefs about their chances of pivotality across elections — whether they believe some elections are close whereas they have no chance of swaying others, (2) how voters may strategize in misreporting their preferences, and most importantly, (3) how voters' preference intensities relate to both their *ability* and their *desire* to vote in a given election. A more holistic mechanism design approach here would build on the wealth of empirical evidence on these topics [Downs, 1957, Feddersen and Pesendorfer, 1996].

6.1.1 Other kinds of trade-offs. In a multi-issue mechanism, voters' stakes are revealed in how they make trade-offs between voting power in election ℓ with voting power in a different election. However, one could also imagine designing a mechanism in which voters must trade off voting power in election ℓ with *some other resource*. One natural option, from a mechanism design perspective, would be currency (or something with similar external economic value). This is precisely the proposal of quadratic voting, which in its theoretical conception requires voters to purchase votes with *money*, thereby revealing voters' stakes in how many votes they buy [Posner and Weyl, 2015]. However, we emphasize that such monetary mechanisms may not recover stakes at all — to do so, they require voters convert their value for voting power to their value for money *at the same rate*, when in reality, less wealthy voters may have higher marginal value for money, in which case they will purchase fewer votes and their stakes will be underestimated.

6.2 Stakes in liquid and deliberative democracy

Our model of *stakes functions* applies in any voting model with latent utilities, and while we only formally define the reweighting of *ranking-based* ballots, our model can easily be extended to study the impact of stakes-dependent reweighting under other ballot formats or decision mechanisms. One extension where studying stakes may especially be of interest is *liquid democracy* [Gölz et al., 2021], where each voter receives one unit of voting power but may delegate it to others. Conceivably, voters might naturally delegate their votes to others in ways that at least partially depend on stakes; in the extreme, one could imagine a pro-social society in which voters delegate their votes to those with the highest stakes, out of a belief that they deserve the most influence. In this case, we might not need to reweight votes explicitly, motivating an interesting question: *if voters delegate their votes in a way that accounts for others' stakes, does the outcome have better distortion?*

One can also conceive of accounting for stakes in ways that go beyond asking who receives the most *voting* power. For example, there is a growing body of work in computer science studying *sortition*, the random process of selecting a "representative" body of citizens to deliberate and make a collective policy recommendation on a given issue [Flanigan et al., 2021]. Although in practice "representation" is typically designed to be proportional to population composition, one could instead implement progressive representation, where representation targets for population groups are at least in part determined by their stakes. In fact, *this is already being tried*: in a deliberative poll in Australia on how to facilitate reconciliation between Indigenous and non-Indigenous groups, Indigenous people — a very small fraction of the overall population, but affected by this decision to an outsize degree — were intentionally over-represented in some deliberation groups [Karpowitz and Raphael, 2016]. More generally, on the topic of representation in the deliberative context, the idea of accounting for the interests of highly affected groups in how representation is decided is already being discussed and advocated [Karpowitz and Raphael, 2016].

REFERENCES

B. Abramowitz, E. Anshelevich, and W. Zhu. 2019. Awareness of voter passion greatly improves the distortion of metric social choice. In *Proceedings of the 15th Conference on Web and Internet Economics (WINE)*. 3–16.

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.

J. J. Andersen, J. H. Fiva, and G. J. Natvik. 2014. Voting when the stakes are high. *Journal of Public Economics* 110 (2014), 157–166.

E. Anshelevich, O. Bhardwaj, E. Elkind, J. Postl, and P. Skowron. 2018. Approximating Optimal Social Choice under Metric Preferences. *Artificial Intelligence* 264 (2018), 27–51.

E. Anshelevich, A. Filos-Ratsikas, N. Shah, and A. A. Voudouris. 2021. Distortion in Social Choice Problems: The First 15 Years and Beyond. arXiv:2103.00911.

Yaron Azrieli. 2018. The price of 'one person, one vote'. *Social Choice and Welfare* 50 (2018), 353–385.

Edana Beauvais and Mark E Warren. 2019. What can deliberative mini-publics contribute to democratic systems? *European Journal of Political Research* 58, 3 (2019), 893–914.

Craig Boutilier, Ioannis Caragiannis, Simi Haber, Tyler Lu, Ariel D Procaccia, and Or Sheffet. 2012. Optimal social choice functions: A utilitarian view. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. 197–214.

C. Boutilier, I. Caragiannis, S. Haber, T. Lu, A. D. Procaccia, and O. Sheffet. 2015. Optimal Social Choice Functions: A Utilitarian View. *Artificial Intelligence* 227 (2015), 190–213.

H. Brighouse and M. Fleurbaey. 2010. Democracy and proportionality. *Journal of Political Philosophy* 18, 2 (2010), 137–155.

Andrew Brinker. 2023. In Cambridge, a battle over affordable housing revives longstanding political tensions. https://www.bostonglobe.com/2023/10/15/business/affordable-housing-cambridge/. Accessed: 2024-02-12.

I. Caragiannis, S. Nath, A. D. Procaccia, and N. Shah. 2017. Subset selection via implicit utilitarian voting. *Journal of Artificial Intelligence Research* 58 (2017), 123–152.

I. Caragiannis and A. D. Procaccia. 2011. Voting almost maximizes social welfare despite limited communication. *Artificial Intelligence* 175, 9-10 (2011), 1655–1671.

A. Casella. 2005. Storable votes. *Games and Economic Behavior* 51, 2 (2005), 391–419.

Anthony Downs. 1957. An economic theory of democracy. *Harper and Row* 28 (1957).

S. Ebadian, A. Kahng, D. Peters, and N. Shah. 2022. Optimized distortion and proportional fairness in voting. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC)*. 563–600.

Timothy J Feddersen and Wolfgang Pesendorfer. 1996. The swing voter's curse. *The American economic review* (1996), 408–424.

B. Flanigan, P. Gölz, A. Gupta, B. Hennig, and A. D. Procaccia. 2021. Fair algorithms for selecting citizens' assemblies. *Nature* 596, 7873 (2021), 548–552.

B. Flanigan, Ariel D. Procaccia, and S. Wang. 2023. Distortion Under Public-Spirited Voting. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*.

Marc Fleurbaey. 2008. Weighted majority and democratic theory. Manuscript.

Archon Fung. 2013. The Principle of Affected Interests: An Interpretation and Defense. *Representation: Elections and beyond* (2013), 236.

William V Gehrlein. 2002. Condorcet's paradox and the likelihood of its occurrence: different perspectives on balanced preferences. *Theory and decision* 52 (2002), 171–199.

P. Gölz, A. Kahng, S. Mackenzie, and A. D. Procaccia. 2021. The fluid mechanics of liquid democracy. *ACM Transactions on Economics and Computation* 9, 4 (2021), 1–39.

A. Kahng and G. Kehne. 2022. Worst-Case Voting When the Stakes Are High. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*, Vol. 36. 5100–5107.

Christopher F Karpowitz and Chad Raphael. 2016. Ideals of inclusion in deliberation. *Journal of Deliberative Democracy* 12, 2 (2016).

Sofia Näsström and Sara Kalm. 2015. A democratic critique of precarity. *Global Discourse* 5, 4 (2015), 556–573.

E. A Posner and E. G. Weyl. 2015. Voting squared: Quadratic voting in democratic politics. *Vanderbilt Law Review* 68 (2015), 441–500.

A. D. Procaccia and J. S. Rosenschein. 2006. The Distortion of Cardinal Preferences in Voting. In *Proceedings of the 10th International Workshop on Cooperative Information Agents (CIA)*. 317–331.

Mark E Warren. 2017. A problem-based approach to democratic theory. *American Political Science Review* 111, 1 (2017), 39–53.

Lirong Xia. 2020. The smoothed possibility of social choice. *Advances in Neural Information Processing Systems* 33 (2020), 11044–11055.

A SUPPLEMENTAL MATERIALS FROM SECTION 1.2

A.1 Connections to existing results

We connect our results to those in three papers. The first two study distortion under the s-unit stakes assumption, and the third assumes utility queries.

Caragiannis et al. [2017] (s-unit stakes, deterministic rules)
This paper assumes sum-unit-stakes. Although this paper proves distortion bounds for both deterministic and randomized rules, we do not discuss their analysis of randomized rules, as such bounds are more directly addressed in later work, described next. Another similarity between our bounds: β_f that our bounds depend on is similar to the dependency of their analysis on alternatives' plurality score.

Upper bounds (deterministic rules): Theorem 1 of their paper proves an $O(m^2)$ upper bound on the distortion of Plurality under sum-unit-stakes (i.e., unit-sum utilities). We can recover this bound via our Theorem 3.4: First observe that κ-upper(sum) $= m$ and κ-lower(sum) $= 1$ (given by utility vectors $\mathbf{1}$ and $\mathbf{1}_1\mathbf{0}_{m-1}$, respectively). Recall also that $\beta_{\text{Plurality}} = 1/m$. By Theorem 3.4, it follows that $\text{dist}^{\text{sum}}(\text{Plurality}) \leq m^2$.

Lower bounds (deterministic rules): Theorem 1 of their paper proves an $\Omega(m^2)$ lower bound on the distortion of any voting rule under sum-unit stakes (unit-sum utilities). We do show a lower-bound on the distortion of all deterministic voting rules, but due to its is general across *any stakes function* (not just sum), our (tight) lower bound is $\Omega(m)$ (Theorem 3.1). However, we can recover the $\Omega(m^2)$ bound specifically for s = sum for *most* voting rules by combining two of our results. First, by Appendix B.6, many voting rules have unbounded distortion under s-unit stakes with respect to *any* s, including sum. Among the remaining rules with $\beta_f > 0$, we can recover a lower bound of $\Omega(m^2)$ (with tighter constants) for Plurality via Proposition B.3: We have κ-upper(sum) from above, and $\tilde{\kappa}$-lower(sum) $= 2$, given by $\mathbf{1}_2\mathbf{0}_{m-2}$. Then, by Proposition B.3, $\text{dist}^{\text{sum}}(\text{Plurality}) \geq (m-1)m/2$. Our bounds here are tighter, improving upon the gap from a factor of 8 to a factor of 2.

Ebadian et al. [2022] (s-unit-stakes, randomized rules)
This paper studies only randomized rules, under both the sum-unit stakes assumption and the max-unit stakes assumption (which they call "range").

Upper bounds (randomized rules): As discussed in Section 3.2, we use our reduction in Appendix B.8 to directly apply their upper bounds on the distortion of Stable Lottery under the sum- and max-unit stakes assumptions (their Theorem 3.4) to prove our upper bound in Theorem 3.11.

Lower bounds (randomized rules): In Theorem 3.7 of their paper, they show a lower bound of $\Omega(\sqrt{m})$ on the distortion of any randomized rule under max-unit-stakes. This complements a previously-known bound by Boutilier et al. [2012] of $\Omega(\sqrt{m})$ on the distortion of any randomized rule under sum-unit stakes. Our lower bound in Theorem 3.11 is weaker than these bounds by a $\log(m)$ factor, but it applies to s-unit stakes for *any 1-homogeneous* s (which includes both sum and max). We suspect our lower bound can be tightened to $\Omega(\sqrt{m})$, which would make it a strict generalization of the existing bound.

Amanatidis et al. [2021] (utility queries, deterministic rules) This paper considers deterministic voting rules with access to one of two kinds of queries: *value* queries, where the voting mechanism can directly ask agents about any one of their utilities; and *comparison* queries, where the voting mechanism can ask agents: "for alternatives a and b, is your utility for a at least d times your utility for b?" Stakes information according to an arbitrary s can be recovered by some number of either type of these queries (trivially, m value queries, but in many cases, far less). Determining the optimal set of queries of these types to recover a given stakes function is outside the scope of this appendix. Thus, when thinking about upper bounds, we will restrict our consideration here to the stakes functions max, which can be recovered by 1 value query. Similar reasoning applies for range. Finally, we remark that their permission of *noisy* queries (i.e., queries within a constant factor of the truth) are related to our robustness results in Theorem 4.1, though due to the generality of our class of stakes functions (and the fact that errors are occurring on *stakes functions' output* rather than utilities directly) requires us to handle additional technicalities.

Upper bounds (deterministic rules): Their upper bound in Theorem 1 shows that their mechanism 1-PRV — equivalent to Plurality under stakes-proportionality with respect to max — gives distortion $O(m)$. This result corresponds to our upper bound on the distortion of plurality under max-proportionality, proven via Theorem 3.4.

Lower bounds (deterministic rules): Their Theorem 7 shows that any single-value query can enable at best $\Omega(m)$ distortion. Our lower bound in Theorem 3.1 generalizes this lower bound, showing that *any system of queries yielding the value of a scalar-valued stakes function*, when paired with a deterministic voting rule, can achieve at best $\Omega(m)$ distortion.[9] Of course, this is not to say that we generalize *all* their lower bounds — they prove several other lower bounds for their setting, which are incomparable to ours.

[9]A necessary step in showing that our lower bound subsumes theirs is arguing that our lower bound actually applies to *any stakes-dependent electoral recomposition*, not just proportional recomposition, which we do in the body of the paper.

B SUPPLEMENTAL MATERIALS FROM SECTION 3

B.1 All deterministic rules have unbounded distortion

Proposition B.1. *For all deterministic rules f, dist$(f) = \infty$.*

Proof. Fix an arbitrary deterministic voting rule f. Consider an election with n voters and $m = 2$ alternatives. For $\epsilon > 0$, let $n/2$ voters have utility vector $(\epsilon, 0)$ and let the other half have $(0, 1)$. Then, half of voters will vote for a and the other half for b. In this example, a or b are indistinguishable to f; suppose it chooses a. Then, the distortion in this instance is $\frac{n/2}{\epsilon n/2} \to \infty$ as $\epsilon \to 0$. □

B.2 Proof of Theorem 3.1

Theorem 3.1 (**lower bound**). *For all s and deterministic f,*

$$\mathsf{dist}^{\mathsf{s}}(f) \geq m - 1.$$

Proof. We will define two instances, U and U', and show that all f must have at least $m - 1$ distortion in one of these two instances. We will construct U, U' in the following way: first, set aside one alternative a', and let the remaining alternatives be $A_\ell = \{a_j | j \in [m] \setminus \{\ell\}\}$. For all ℓ, when we write A_ℓ in a ranking it represents a ranking over all the alternatives within it, *in increasing order of index.* Divide voters in into $m - 1$ groups, and consider a voter i in group ℓ: we will assign utility vectors to these voters so that their ranking $\pi_i = a_\ell \succ a' \succ A_\ell$. We display i's utility vectors $\mathbf{u}_i$ and $\mathbf{u}'_i$, as given by U and U' respectively, in *sorted* order, to emphasize how their utilities correspond to their resulting ranking:

alternative:	a_ℓ	$\succ$	a'	$\succ$	A_ℓ
sorted $\mathbf{u}_i$ for $i \in$ group ℓ:	1		1		$0 \dots 0$
sorted $\mathbf{u}'_i$ for $i \in$ group ℓ:	1		0		$0 \dots 0$

We now make three observations:

(1) $\mathsf{hist}(U) \equiv \mathsf{hist}(U')$ — that is, the utility matrices induce the same preference histogram. This is true because for every ℓ, voters in the ℓ-th group of U and U' have the same ranking.
(2) $\mathsf{hist}^s(U) \equiv \mathsf{hist}(U)$ and $\mathsf{hist}^s(U') \equiv \mathsf{hist}(U')$ — that is, the s-proportional profiles are identical to the standard profiles for both utility matrices. This is because within each utility matrix, all voters have the same ordered utility vector and thus have the same stakes.
(3) $\mathsf{sw}(a', U) = n$ while $\mathsf{sw}(a', U') = 0$. Moreover, $\mathsf{sw}(a_\ell, U) = \mathsf{sw}(a_\ell, U') = n/(m-1)$ for all $\ell \in [m-1]$.

We distinguish between two cases, depending on whether $f(\mathsf{hist}(U)) = a'$ or $f(\mathsf{hist}(U)) \neq a'$. If $f(\mathsf{hist}(U)) = a'$, by (1), we also have that $f(\mathsf{hist}(U')) = a'$. Then, since $\mathsf{sw}(a', U') = 0$,

$$\mathsf{dist}^{s}_{U'}(f) \overset{(2)}{=} \mathsf{dist}_{U'}(f) = \frac{\mathsf{sw}(a_1, U')}{\mathsf{sw}(a', U')} \overset{(3)}{=} \frac{n/(m-1)}{0} = \infty.$$

If $f(\mathsf{hist}(U)) \neq a'$, then there must exist some $\ell \in [m-1]$ such that $f(\mathsf{hist}(\mathbf{u})) = a_\ell$. Then, fixing this ℓ,

$$\mathsf{dist}^{\mathsf{s}}_{U}(f) \overset{(2)}{=} \mathsf{dist}_{U}(f) = \frac{\mathsf{sw}(a', U)}{\mathsf{sw}(a_\ell, U)} \overset{(3)}{=} \frac{1}{1/(m-1)} = m - 1.$$ □

B.3 Theorem 3.4 holds when κ's are defined with range instead of max

Observation B.2. *The bound in Theorem 3.4 remains true also for a slightly different definition of the coefficients* κ-lower(s), κ-upper(s) *where* max$(\cdot)$ *is replaced by* range$(\cdot)$*,*

$$\kappa\text{-upper}(s) := \sup_{\mathbf{u} \in \mathbb{R}^m_{\geq 0}} \frac{s(\mathbf{u})}{\mathit{range}(\mathbf{u})}, \quad \textit{and} \quad \kappa\text{-lower}(s) := \inf_{\mathbf{u} \in \mathbb{R}^m_{\geq 0}} \frac{s(\mathbf{u})}{\mathit{range}(\mathbf{u})}.$$

Proof. Let $U \in \mathbb{R}_{\geq 0}^{m \times n}$ be any utility matrix. Then, let $\tilde{U}$ denote the utility matrix in which each agent $i's$ utility vector $\mathbf{u}_i$ is altered by

$$\tilde{u}_i(a) = u_i(a) - \min_{a \in [m]} u_i(a),$$

i.e., the utilities are shifted down such that each voter's minimum utility is 0. Then, letting $c := \sum_{i \in [N]} \min_a u_i(a)$, we obtain that

$$\frac{\mathsf{sw}(a^*, U)}{\mathsf{sw}(a', U)} \leq \frac{\mathsf{sw}(a^*, U) - c}{\mathsf{sw}(a', U) - c} = \frac{\mathsf{sw}(a^*, \tilde{U})}{\mathsf{sw}(a, \tilde{U})}.$$

Then, we may restrict the arguments in the proof of Theorem 3.4 to utility vectors with zero minimum entry. This leads to a bound where we may use, instead of κ-upper(s) and κ-lower(s)

$$\sup_{\mathbf{u} \in \mathbb{R}_{\geq 0}^m : \min_{a \in [m]} u(a) = 0} \frac{s(\mathbf{u})}{\max(\mathbf{u})} \text{ and } \inf_{\mathbf{u} \in \mathbb{R}_{\geq 0}^m : \min_{a \in [m]} u(a) = 0} \frac{s(\mathbf{u})}{\max(\mathbf{u})}$$

in place of κ-upper(s) and κ-lower(s). We may further upper and lower bound these last two quantities, respectively, by

$$\sup_{\mathbf{u} \in \mathbb{R}_{\geq 0}^m : \min_a u(a) = 0} \frac{s(\mathbf{u})}{\max(\mathbf{u})} = \sup_{\mathbf{u} \in \mathbb{R}_{\geq 0}^m : \min_a u(a) = 0} \frac{s(\mathbf{u})}{\mathrm{range}(\mathbf{u})} \leq \sup_{\mathbf{u} \in \mathbb{R}_{\geq 0}^m} \frac{s(\mathbf{u})}{\mathrm{range}(\mathbf{u})},$$

$$\inf_{\mathbf{u} \in \mathbb{R}_{\geq 0}^m : \min_a u(a) = 0} \frac{s(\mathbf{u})}{\max(\mathbf{u})} = \inf_{\mathbf{u} \in \mathbb{R}_{\geq 0}^m : \min_a u(a) = 0} \frac{s(\mathbf{u})}{\mathrm{range}(\mathbf{u})} \geq \inf_{\mathbf{u} \in \mathbb{R}_{\geq 0}^m} \frac{s(\mathbf{u})}{\mathrm{range}(\mathbf{u})},$$

and we then in particular obtain a distortion upper bound with the two expressions on the right hand side in place of κ-upper(s) and κ-lower(s). □

B.4 Proof of Lemma 3.6

Proof. Fix any deterministic voting rule f, and define the quantity

$$\kappa_f = \min_{\mathbf{h} \in \Delta(S_m)} \min_{a \neq f(\mathbf{h})} \sum_{\pi \in S_m} h_\pi \mathbb{I}(f(\mathbf{h}) \succ_\pi a),$$

which captures the minimum fraction of people by whom the winner $f(\boldsymbol{\pi})$ ranked ahead of any other given alternative a. In [Flanigan et al., 2023], it is shown that for any voting rule f, we have that

$$\kappa_f \leq \kappa_{\textsc{Minimax}} = 1/m,$$

where Minimax is the voting rule which chooses the alternative a that suffers the least severe worst pairwise defeat; see [Flanigan et al., 2023] for details. Moreover, we have that for any histogram profile $\mathbf{h}$ and any alternative $a \neq f(\mathbf{h})$,

$$\sum_{\pi \in S_m} h_\pi \mathbb{I}(\pi^{-1}(f(\mathbf{h})) = 1) \leq \sum_{\pi \in S_m} h_\pi \mathbb{I}(f(\mathbf{h}) \succ_\pi a)$$

It follows that $\beta_f \leq \kappa_f \leq 1/m$, which proves the first part of the claim.

Now, for the second part of the claim: the fact that $\beta_{\textsc{Plurality}} \geq 1/m$ follows immediately its definition: there always exists an alternative which is first-ranked in at least a $1/m$ fraction of the population – therefore, the Plurality winner also has to rank first at least in a $1/m$ fraction of the population. □

B.5 Proof of Proposition B.3

Proposition B.3. *For all s,*

$$\mathrm{dist}^{\mathrm{s}}(\textsc{Plurality}) \geq (m-1)\cdot \kappa\text{-upper}(s)/\tilde{\kappa}\text{-lower}(s).$$

Proof. Formally, we define $\tilde{\kappa}^{\mathrm{lower}}$ as

$$\tilde{\kappa}^{\mathrm{lower}} = \inf_{u\in\mathcal{U}} \frac{s(u)}{\max \mathbf{u}}, \qquad \mathcal{U} := \left\{u \in \mathbb{R}^m_{\geq 0} : u_1 = u_2 \geq \cdots \geq u_m = 0\right\}. \tag{4}$$

We will construct an instance which exhibits distortion of the desired order.

Step 1: Designing the ordered utilities. There are two population groups: one *high-stake* population group which we call G_1 and on *low-stake* population group which we call G_2. We denote the proportional group size of G_1 by $p = |G_1|/n \in (0,1)$, $1-p = |G_2|/n$. The exact value of p will be determined later in Step 3 of this proof.

Since we are considering proportional recomposition, we may assume without loss of generality that across agents, their maximal utility is equal to 1. Suppose that u^{upper} is an ordered utility vector which maximizes the supremum in κ-upper, such that $\max_{a\in[m]} u^{\mathrm{upper}}(a) = 1$. Similarly, let u^{lower} denote the utility vector in $\mathcal{U}$ that minimizes the infimum in (4). Now, we assign to G_1 the ordered utility vector u^{upper}, and to G_2 the ordered utility vector u^{lower}. Then, agents in these two population groups have respective stakes of

$$s(u^{\mathrm{upper}}) = \kappa\text{-upper}, \qquad s(u^{\mathrm{lower}}) = \tilde{\kappa}^{\mathrm{lower}}.$$

Step 2: Designing the rankings.

- In group G_1, we first-rank an alternative a' – this alternative, by appropriate choice of p, will later turn out to be the winner of the plurality election. The second to last ranked alternatives in group G_1 can be chosen arbitrarily.
- In group G_2, the first-rank positions are divided up equally between the remaining $m-1$ alternatives in $[m]\setminus\{a'\}$. Out of those $m-1$ alternatives, we choose an arbitrary alternative which we will make the highest-welfare alternative, called a^*. This alternative a^* is ranked second throughout the group G_2, whenever it does not rank first.
- Finally, we also specify that the alternative a' is ranked last throughout group G_2. The remaining places in G_2's preference profile may be filled arbitrarily.

Step 3: Specifying the group size p. It remains to calculate p. Since G_1 has stakes κ-upper and G_2 has stakes $\tilde{\kappa}^{\mathrm{lower}}$, the stakes-weighted plurality score obtained by a' is $p\kappa$-upper. Any other alternative $a \neq a'$ obtains a stakes-weighted plurality score of $(1-p)\tilde{\kappa}^{\mathrm{lower}}/(m-1)$. Thus, a' winning the election amounts to the inequality

$$p\kappa\text{-upper} \geq \frac{1-p}{m-1}\tilde{\kappa}^{\mathrm{lower}} \iff p(\kappa\text{-upper}+\frac{\tilde{\kappa}^{\mathrm{lower}}}{m-1}) \geq \frac{\tilde{\kappa}^{\mathrm{lower}}}{m-1} \iff p \geq \frac{\tilde{\kappa}^{\mathrm{lower}}}{\tilde{\kappa}^{\mathrm{lower}} + (m-1)\kappa\text{-upper}}.$$

Thus, let us set p to be equal to the last expression, i.e.

$$p = \frac{|G_1|}{n} = \frac{\tilde{\kappa}^{\mathrm{lower}}}{\tilde{\kappa}^{\mathrm{lower}} + (m-1)\kappa\text{-upper}}.$$

With this choice of p, we notice that

$$\frac{\mathrm{sw}(a', U)}{n} = p, \qquad \text{and} \qquad \frac{\mathrm{sw}(a^*, U)}{n} \geq \frac{1}{n}\sum_{i\in G_2} u_i(a^*) = 1-p,$$

since agents in G_2 have utility 1 for a^*, and agents in G_1 *may* have positive utility for a^*. In conclusion, the distortion in this instance is lower bounded by

$$\frac{\mathrm{sw}(a^*,U)}{\mathrm{sw}(a',U)} \geq \frac{1-p}{p} = \frac{\frac{(m-1)\kappa\text{-upper}}{\tilde{\kappa}^{\text{lower}}+(m-1)\kappa\text{-upper}}}{\frac{\tilde{\kappa}^{\text{lower}}}{\tilde{\kappa}^{\text{lower}}+(m-1)\kappa\text{-upper}}} = \frac{(m-1)\kappa\text{-upper}}{\tilde{\kappa}^{\text{lower}}}.$$

□

B.6 Proof: $\beta_f = 0$ for many established voting rules

Observation B.4. *$\beta_f = 0$ for many established voting rules.*

This observation is shown via a simple instance. Before presenting this instance, we define the voting rules we will address.

Voting rules. Borda Count and Veto are positional scoring rules, which are rules defined by a scoring vector $w \in [0,1]^m$ with j-th entry w_j. In these scoring rules, an alternative receives w_j points for each voter who ranks it j-th, and the winner in a given profile is the alternative with the most points. Borda Count is defined by the linearly-decreasing scoring vector $\mathbf{w} = (1, (m-2)/(m-1), \ldots, 1/(m-1), 0)$, and Veto is defined by $\mathbf{w} = \mathbf{1}_{m-1}\mathbf{0}_1$. We also consider the entire class of *Condorcet-consistent rules.* To define this class, we say that *a pairwise-dominates a'* in $\mathbf{h}$ if a is ranked ahead of a' in at least half of the electorate. We say that $\mathbf{h}$ has a *Condorcet winner a* if a pairwise-dominates all other alternatives. A Condorcet-consistent rule is one which $f(\mathbf{h})$ will be the Condorcet winner on all profiles $\mathbf{h}$ in which a Condorcet winner exists. We will consider this large class of voting rules as a whole, but will not consider any specific rule in this class.

Instance. Indeed, consider the following instance with 4 alternatives, a, b, c, d:

- 1 voter has $c \succ a \succ d \succ b$
- $n/3 - 1$ voters have $c \succ a \succ b \succ d$
- $n/3$ voters have $b \succ a \succ c \succ d$
- $n/3$ voters have $d \succ a \succ b \succ c$

Then, a is ranked ahead of any other alternative by $2/3$ of voters, and is the Condorcet winner; it will also be the Borda winner, and the Veto winner. Yet, it is never ranked first.

B.7 Proof of Theorem 3.11

Theorem 3.11 (**lower bound**). *For all 1-homogeneous s, randomized f,* $\mathrm{dist}^s(f) \geq \frac{\sqrt{m}}{10+3\log m}$.

Proof. Define the vector $\mathbf{1}_z\mathbf{0}_{z'}$ to be the vector consisting of z ones followed by z' zeroes.

Case 1: Suppose that there exists some $z \leq (\log m) - 1$ such that $s(\mathbf{1}_{z+1}\mathbf{0}_{m-z-1})/s(\mathbf{1}_z\mathbf{0}_{m-z}) \leq e$. Fix this z. We now design a utility instance and associated preference histogram which exhibits a distortion of the order $\sqrt{m/\log m}$.

Step 1: Designing the rankings. We begin by designing the preference histogram. We divide the population into $m/\log m$ groups

$$G_1, \ldots G_{m/\log m}.$$

Let alternatives $1, \ldots, m/\log m$ occupy the first positions in each of the groups $G_1, \ldots G_{m/\log m}$, respectively. Similarly, we occupy the second to z-th rank of those groups by following alternatives:

$$\begin{array}{rcccc} \text{Rank:} & 1 & 2 & \ldots & z \\ \text{Group } G_1\text{:} & 1 & m/\log m + 1 & \ldots & (z-1)m/\log m + 1 \\ & \vdots & & & \vdots \\ \text{Group } G_{m/\log m}\text{:} & m/\log m & 2m/\log m & \ldots & zm/\log m. \end{array}$$

Next, we also divide the population into $\sqrt{m}$ parts $H_1, ..., H_{\sqrt{m}}$ of equal size, based on which alternatives occupy the $(z+1)$-the position. We may design this partition in a way such that

$$\forall k \in [\sqrt{m}] : \left|\{l \in [m/\log m] : H_k \cap G_l \neq \emptyset\}\right| \leq \frac{\sqrt{m}}{\log m} + 2.$$

Intuitively, this is because the groups H_k are *larger* by a factor of $\sqrt{m}/\log m$ than the groups G_l. We may thus pick the partition into H_k such that each H_k overlaps with at most $\sqrt{m}/\log m + 2$ many groups G_l. For each $k \in [\sqrt{m}]$, we assign the $(z+1)$-th position in group H_k to be occupied by the alternative $zm/\log m + k$. Finally, we fill the rest of the positions in the preference histogram – i.e. the $(z+2)$-th to last ranks – arbitrarily.

Step 2: Designing the utilities. Amongst the $\sqrt{m}$ alternatives which are ranked in the $(z+1)$-th position, there must exist one alternative which we call $\bar{a}$ which is chosen by the voting rule f with probability at most $1/\sqrt{m}$. That is, if $\mathbf{h}$ denotes the preference histogram constructed in Step 1, then

$$f_{\bar{a}}(\mathbf{h}) \leq 1/\sqrt{m}.$$

Let $H_{\bar{k}}$ be the unique group which ranks $\bar{a}$ in the $(z+1)$-th position. Now, we assign utilities as follows. Define the following ratio of stakes:

$$c_z := \frac{s(\mathbf{1}_{z+1}\mathbf{0}_{m-z-1})}{s(\mathbf{1}_z\mathbf{0}_{m-z})} \leq e.$$

- **Group $H_{\bar{k}}$.** We assign to agents in H_k the ranked utilities $s(\mathbf{1}_{z+1}\mathbf{0}_{m-z-1})$.
- **Remainder.** In the remaining population $H_{\bar{k}}^c$, we assign the ranked utilities $c_z \cdot s(\mathbf{1}_z\mathbf{0}_{m-z})$.

These ordered utilities, together with the rankings designed in Step 1, determine a utility matrix which we call U.

(1) The alternative $\bar{a}$ has average utility $\mathsf{sw}(\bar{a}, U) = 1/\sqrt{m}$.
(2) All other alternatives $a \neq \bar{a}$ have average utility at most $\mathsf{sw}(a, U) = c_z \log m/m \leq e \log m/m$.
(3) By the homogeneity of the stakes function $s(\cdot)$, all voters have equal stakes. Therefore, we have that $\mathsf{hist}^s(U) = \mathsf{hist}(U) = \mathbf{h}$, and thus also

$$f(\mathsf{hist}^s(U)) = f(\mathbf{h}).$$

In particular, $\bar{a}$ is chosen by the voting rule with probability at most $1/\sqrt{m}$ in $f(\mathsf{hist}^s(U))$.

Together, these observations yield that

$$\mathbb{E}\left[\mathsf{sw}(f(\mathsf{hist}^s(U)))\right] \leq \frac{e \log m}{m} + \frac{1}{\sqrt{m}}\frac{1}{\sqrt{m}} = \frac{e \log m + 1}{m},$$

and thus the f in Case 1 is at least

$$\frac{\max_a \mathsf{sw}(a, U)}{\mathbb{E}\left[\mathsf{sw}(f(\mathsf{hist}^s(U)))\right]} \geq \frac{1/\sqrt{m}}{(1 + e \log m)/m} = \frac{\sqrt{m}}{1 + e \log m}.$$

Case 2: It remains to treat the case when the premise of Case 1 is not fulfilled, that is, for every $z \leq \log m - 1$, it holds that $s(\mathbf{1}_{z+1}\mathbf{0}_{m-z-1})/s(\mathbf{1}_z\mathbf{0}_{m-z}) \geq e$. By multiplying this equality for all $z = 2, \ldots, \log m - 1$, it follows that

$$\frac{s\left(\mathbf{1}_{\log(m)-1}\mathbf{0}_{m-\log(m)+1}\right)}{s(\mathbf{1}_1\mathbf{0}_{m-1})} \geq 2^{\log m - 2} \geq \frac{m}{e^2}. \tag{5}$$

Now let us consider a histogram profile where the population is divided in $\sqrt{m}$ many equal sizes groups, which first-rank alternatives $1, ...\sqrt{m}$, respectively. We fill up the remaining positions in the histogram arbitrarily. Denote this histogram by $\mathbf{h}$.

We now assign utilities to induce $\mathbf{h}$. There must exist one alternative among the $\sqrt{m}$ first-ranked alternatives that receives $\leq 1/\sqrt{m}$ probability of selection by $f(\mathbf{h})$. Let us call this alternative a^*, and let us call the group which ranks a^* first G.

- **Group** G**.** In this group, we assign the ordered utility vector $\mathbf{1}_1\mathbf{0}_{m-1}$.
- **Group** G^c**.** In the remainder of the population, we assign the ordered utility vector

$$\frac{s(\mathbf{1}_1\mathbf{0}_{m-1})}{s(\mathbf{1}_{\log(m)-1}\mathbf{0}_{m-\log(m)+1})} \cdot \mathbf{1}_{\log(m)-1}\mathbf{0}_{m-\log(m)+1}$$

Let us denote the resulting utility matrix by U. We observe the following.

(1) The average utility of a^* is at least $\mathsf{sw}(a^*, U)/n \geq 1/\sqrt{m}$.
(2) By equation (5), the average utility of any other alternative $a \neq a^*$ is at most

$$\frac{\mathsf{sw}(a, U)}{n} \leq \frac{e^2}{m}.$$

(3) All voters have equal stakes. Therefore $f(\mathbf{h}) = f(\mathsf{hist}(U)) = f(\mathsf{hist}^s(U))$ and we may estimate

$$\mathbb{E}[\mathsf{sw}(f(\mathsf{hist}^s(U)), U)] \leq \frac{1}{\sqrt{m}}\frac{1}{\sqrt{m}} + \frac{e^2}{m^2} \leq \frac{10}{m}.$$

We obtain an overall distortion of at least

$$\mathsf{dist}^s_U(f) \geq \frac{\sqrt{m}}{10},$$

and the proof is complete. □

B.8 Proof of Lemma 3.13

B.8.1 Theorem B.5: reduction for rational-valued histograms. Here, we state and prove the reduction assuming $\mathsf{hist}^s(U)$ has only rational entries, which we ensure by restricting to rational utility matrices $U \in \mathbb{Q}^{n\times m}_{\geq 0}$ and *rationality-preserving* stakes functions s (i.e., $s(\mathbf{u}) \in \mathbb{Q}$ whenever $\mathbf{u} \in \mathbb{Q}^m_{\geq 0}$).

Theorem B.5. *Let f be a voting rule, s a rationality-preserving and 1-homogeneous stakes function, and let $\mathcal{U}_s$ be the set of all rational utility matrices satisfying the s-unit-stakes assumption. Then,*

$$\sup_{n\geq 1} \sup_{U\in\mathcal{U}_s} dist_U(f) = \sup_{n\geq 1} \sup_{U\in\mathbb{Q}^{n\times m}_{\geq 0}} dist^s_U(f).$$

Proof. We show the claimed equality by separately proving the directions '$\leq$' and '$\geq$'. In order to see the direction '$\leq$', we note that for any unit-stakes utility matrix $U \in \mathcal{U}_s$, $\mathsf{hist}(U) = \mathsf{hist}^s(U)$: the standard and stakes-proportional histograms are the same. Therefore, $\mathsf{dist}_U(f) = \mathsf{dist}^s_U(f)$. Taking suprema over $n \geq 1$ and $U \in \mathcal{U}_s$, we obtain the '$\leq$' direction.

It remains to show '$\geq$'. In order to prove this direction, we fix any utility matrix $U \in \mathbb{Q}^{n\times m}_{\geq 0}$, and construct a unit-stakes utility matrix $\tilde{U}$ such that $\mathsf{dist}_{\tilde{U}}(f) = \mathsf{dist}^s_U(f)$. We let

$$\bar{s}_i = \frac{s(\mathbf{u}_i)}{\sum_{i\in[n]} s(\mathbf{u}_i)}, \quad i \in [n]$$

be the weights with which voter i is represented in the stakes-recomposed election. Since $\bar{s}_i \in \mathbb{Q}$, there exists some $\tilde{n}$ such that $\bar{s}_i\tilde{n}$ is again an integer for each $i \in [n]$. We fix such an $\tilde{n}$ and now construct a utility matrix $\tilde{U} \in \mathbb{Q}^{\tilde{n}\times m}_{\geq 0}$ for which f (without taking into account stakes) exhibits the same distortion as U (while accounting for stakes).

- We divide the electorate of $\tilde{n}$ into n groups, each of them of size $\bar{s}_i\tilde{n}$. Call these groups $G_1, ...G_n$.
- Within each group G_i, voters have the same ranking $\pi_i(U)$ as voter i in U. However, they possess *scaled* utilities $\mathbf{u}_i/s(\mathbf{u}_i)$.

Then we notice that by definition, $\mathsf{hist}(\tilde{U}) = \mathsf{hist}^s(U)$, and therefore also $f(\mathsf{hist}(\tilde{U})) = f(\mathsf{hist}^s(U))$. Moreover, since s is 1-homogeneous, it holds that for all i,

$$s\Big(\frac{\mathbf{u}_i}{s(\mathbf{u}_i)}\Big) = \frac{1}{s(\mathbf{u}_i)} s(\mathbf{u}_i) = 1,$$

which yields that $\mathcal{U}_s$ satisfies the unit-stakes property. Moreover, for all alternatives $a \in [m]$, it holds that

$$\frac{\mathsf{sw}(a, U)}{n} = \sum_{i \in [n]} u_i(a) = \frac{\sum_i s(\mathbf{u}_i)}{n} \sum_{i \in [n]} \bar{s}_i \frac{u_i(a)}{s(\mathbf{u}_i)} = \sum_{i \in [n]} s(\mathbf{u}_i) \cdot \frac{\mathsf{sw}(a, \tilde{U})}{n} = \sum_{i \in [n]} s(\mathbf{u}_i) \frac{\tilde{n}}{n} \cdot \frac{\mathsf{sw}(a, \tilde{U})}{\tilde{n}}.$$

Since $\sum_{i \in [n]} s(\mathbf{u}_i) \frac{\tilde{n}}{n}$ is a fixed constant independent of i and a, it follows that the average utilities in U and $\tilde{U}$ are equal up to multiplication with a fixed constant — thus distortion is preserved. □

B.8.2 Theorem B.7: Extension of Theorem B.5 to real-valued histograms. Under an additional *very mild* restrictions on the voting rule f, it is possible to prove the correspondence between stakes-based procedures and unit-stakes assumptions from Theorem B.5 not just for rational utilities, but for all real-valued utility functions. We term this assumption for f to be *rationally approximable*, which amount to the outcome of $f(\mathbf{h})$ for any preference histogram being *well-approximated* by some preference histogram $\tilde{\mathbf{h}}$ with only rational entries.

Definition B.6 (Rationally approximable rules). We say that a (deterministic or randomized) voting rule $f : \Delta(S_m) \to \Delta([m])$ is 'rationally approximable' if for every $\mathbf{h} \in \Delta(S_m)$ and every $\varepsilon > 0$ there exists another histogram $\tilde{\mathbf{h}} \in \mathbb{Q}_{\geq 0}^{n \times m}$ with only rational entries such that

$$\sup_{\pi \in S_m} |h_\pi - \tilde{h}_\pi| \leq \varepsilon \quad \text{and} \quad \sup_{a \in [m]} |f_a(\mathbf{h}) - f_a(\tilde{\mathbf{h}})| \leq \varepsilon,$$

where $f_a(\mathbf{h})$ denotes the win probability of a in $f(\mathbf{h})$.

Theorem B.7. *For any* 1*-homogeneous stakes function s and any voting rule $f : \Delta([m!]) \to \Delta([m])$, we have that*

$$\sup_{n \geq 1} \sup_{U \in \mathcal{U}_s} \mathit{dist}_U(f) \leq \mathit{dist}^s(f).$$

If additionally s is 1-homogeneous and f is either (i) weakly locally constant or (ii) continuous, then the reverse inequality is also true,

$$\sup_{n \geq 1} \sup_{U \in \mathcal{U}_s} \mathit{dist}_U(f) \geq \mathit{dist}^s(f).$$

Proof of Theorem B.7. The first inequality is immediately implied by the fact that for any $U \in \mathcal{U}_s$, the stakes-recomposed electorate is identical to the original electorate. Indeed, in this case stakes-based election yields the same outcome as the non-stakes-based election, $f(\mathsf{hist}(U)) = f(\mathsf{hist}^s(U))$, so that $\mathsf{dist}_U(f) = \mathsf{dist}_U^s(f)$. It thus only remains to prove the reverse inequality.

Let us fix an arbitrary $n \geq 1$ and utility matrix $U \in \mathbb{R}^{n \times m}$, and let $\mathsf{hist}^s(U) \in \Delta(S_m)$ denote the stakes-recomposed profile corresponding to U. Without loss of generality, we may assume both $\mathsf{sw}(a^*, U) > 0$ (since otherwise $U = 0$) and

$$\mathbb{E}\Big[\mathsf{sw}(f(\mathsf{hist}^s(U)), U)\Big] > 0,$$

since otherwise $\mathsf{dist}_U^s(f) = \infty$ and there remains nothing to prove. By Proposition B.8, given *any* $\rho > 0$ we may choose a unit-stakes utility matrix $\tilde{U} \in \mathbb{R}_{\geq 0}^{\tilde{n} \times m}$ such that

$$\sup_{a \in [m]} |f_a(\mathsf{hist}^s(U)) - f_a(\mathsf{hist}(\tilde{U}))| \leq \rho \qquad \text{and} \qquad \sup_{a \in [m]} \left| \frac{\mathsf{sw}(a, U)}{n} - \frac{\mathsf{sw}(a, \tilde{U})}{\tilde{n}} \right| \leq \rho.$$

These two properties, taken together, imply the convergence

$$\Big|\mathbb{E}\big[\frac{\mathsf{sw}(f(\mathsf{hist}^s(U)),U)}{n}\big] - \mathbb{E}\big[\frac{\mathsf{sw}(f(\mathsf{hist}(\tilde{U})),\tilde{U})}{\tilde{n}}\big]\Big| \xrightarrow{\rho\to 0} 0,$$

as well as the convergence

$$\Big|\max_{a\in[m]}\frac{\mathsf{sw}(a,U)}{n} - \max_{a\in[m]}\frac{\mathsf{sw}(a,\tilde{U})}{\tilde{n}}\Big| \xrightarrow{\rho\to 0} 0.$$

Taken together, this implies that

$$|\mathsf{dist}_U^s(f) - \mathsf{dist}_{\tilde{U}}(f)| \xrightarrow{\rho\to 0} 0,$$

which proves the claim. □

Proposition B.8 (Approximation of social welfares)**.** *Suppose f is a rationally approximable voting rule. Let $U \in \mathbb{R}^{n\times m}$ be any non-zero utility matrix. Then, for any $\rho > 0$ there exists some large enough $\tilde{n}$ and a unit-stakes utility matrix $\tilde{U} \in \mathbb{R}^{\tilde{n}\times m}$ such that*

- *The election outcomes are close,*

$$\sup_{a\in[m]} |f_a(hist^s(U)) - f_a(hist(\tilde{U}))| \leq \rho.$$

- *For all $a \in [m]$, the average utilities in U and $\tilde{U}$ are close,*

$$\Big|\frac{sw(a,U)}{n} - \frac{sw(a,\tilde{U})}{\tilde{n}}\Big| \leq \rho.$$

Proof. Let $\varepsilon > 0$ be arbitrary and fix any U. By Definition B.6, we can choose some $\tilde{\mathbf{h}} \in \mathbb{Q}_{\geq 0}^{S_m}$ with rational coefficients such that

$$\sup_{\pi\in S_m} |\mathsf{hist}_\pi^s(U) - \tilde{h}_\pi| \leq \varepsilon \quad \text{and} \quad \sup_{a\in[m]} |f\big(\mathsf{hist}_\pi^s(U)\big) - f_a(\tilde{h})| \leq \varepsilon,$$

Step 1: Construction of utility matrix which induces $\tilde{\mathbf{h}}$. Since $\tilde{\mathbf{h}} \in \mathbb{Q}_{\geq 0}^{S_m}$ only has rational coefficients, there exists some electorate with $\tilde{n}$ many voters and preferences $(\tilde{\pi}_i : i \leq \tilde{n})$ such that for each $\pi \in S_m$, exactly a $\tilde{h}_\pi$ fraction of the voters have ranking π. Now, we construct a unit-stakes utility matrix $\tilde{U} \in \mathcal{U}_s \cap \mathbb{R}^{\tilde{n}\times m}$ which induces those rankings to the $\tilde{n}$ voters, and which in turn will induce the profile $\tilde{\mathbf{h}}$, $\mathsf{hist}(\tilde{U}) = \tilde{\mathbf{h}}$. To this end, let

$$\bar{s}_i := \frac{s(u_i)}{\sum_{i\in[n]} s(u_i)}, \qquad \sum_{i\in[n]} \bar{s}_i = 1,$$

denote the weights corresponding to each voter i's preferences in the stakes-recomposed electorate. Since s is 1-homogeneous, we may assume without loss of generality that $\sum_{i\in[n]} s(u_i) = n$, by simply scaling the utilities (note that this leaves $\mathsf{hist}^s(U)$ and also $\mathsf{dist}_U^s(f)$ unchanged). We partition in the new 'unit-stakes electorate' (which consists of $\tilde{n}$ voters) into $n+1$ parts, which we denote by $G_1, ..., G_{n+1}$. Within each of those groups, voters share the same ordered utility vector.

Groups $G_1, ...G_n$. The first n groups $G_1, ..., G_n$ are specified as follows. Voters in group i have the utilities $\frac{u_i}{s(u_i)}$, i.e., the same utilities as voter i in the original electorate, but *scaled* to unit-stakes. In particular, voters in group G_i will inherit the same ranking π_i as the $i-th$ voter from the original electorate. Let the (fraction) size of the i-th group be denoted by g_i, i.e., $g_i = |G_i|/\tilde{n}$. We now determine those sizes. Since

$$\sup_{\pi\in S_m} \big|\tilde{\mathbf{h}}_\pi - \mathsf{hist}_\pi^s(U)\big| \leq \varepsilon,$$

we can now choose the $(g_i : i \in n)$ in such a way such that simultaneously, the following properties are satisfied. First, $g_i \in [\bar{s}_i - \varepsilon, \bar{s}_i]$, and second, for every $\pi \in S_m$,

$$\sum_{i \in n} g_i \mathbb{I}(\pi_i = \pi) \leq \tilde{\mathbf{h}}_\pi. \tag{6}$$

The first property states that the group size G_i does not exceed the amount of representation of voter i in the stakes-recomposed electorate $\bar{s}_i$. The second property states that by assigning group sizes g_i, compared to the histogram $\tilde{\mathbf{h}}$, none of the rankings is *overrepresented.* Note that

$$\sum_i g_i \leq \sum_i \bar{s}_i \leq 1, \quad \text{and} \quad \sum_i g_i \geq \sum_i \bar{s}_i - \varepsilon \geq 1 - n\varepsilon.$$

Group G_{n+1}**.** This group constitutes the remainder of the population. Within this group, everyone has the same *ordered utility vector*, but not the same rankings of alternatives. In this group, we assign the ordered utility vector $(x, 0, \ldots, 0)$, where x is given by $x = s((1, 0, \ldots, 0))^{-1} > 0$. Note that x is the (unique) constant such that $s((x, 0, \ldots, 0)) = 1$. In terms of the orderings of alternatives in group G_{n+1}, we assign the exact rankings which are needed to complete the correct histogram $\tilde{\mathbf{h}}$ which we aim to realize. Since from Groups $G_1, \ldots, G_n$, none of the rankings $\pi \in S_m$ was overrepresented compared to $\tilde{\mathbf{h}}$ – see equation (6) – this is possible. The group G_{n+1} has size at most $n\varepsilon$.

Let us denote the utility matrix which arises from this construction by $\tilde{U} \in \mathbb{R}_{\geq 0}^{\tilde{n} \times m}$.

Step 2: Approximation of social welfares. It remains to check that the distortion $\mathsf{dist}_{\tilde{U}}(f)$ induced by $\tilde{U}$ approximates the distortion $\mathsf{dist}_U^s(f)$ for the stakes-based election. To this end, we upper and lower bound the difference in average utilities induced by U and $\tilde{U}$, respectively. First, recalling that $\sum_i s(u_i) = n$, we have the lower bound

$$\begin{aligned}
\frac{\mathsf{sw}(a, \tilde{U})}{\tilde{n}} - \frac{\mathsf{sw}(a, U)}{n} &\geq \sum_{i=1}^n g_i \frac{u_i(a)}{s(u_i)} - \frac{1}{n} \sum_{i=1}^n u_i(a) \\
&\geq \sum_{i=1}^n (\bar{s}_i - \varepsilon) \frac{u_i(a)}{s(u_i)} - \frac{1}{n} \sum_{i=1}^n u_i(a) \\
&\geq \sum_{i=1}^n \frac{s(u_i)}{\sum_{j \in [n]} s(u_j)} \frac{u_i(a)}{s(u_i)} - \frac{1}{n} \sum_{i=1}^n u_i(a) - \sum_{i=1}^n \varepsilon \frac{u_i(a)}{s(u_i)} \\
&= -\varepsilon \sum_{i=1}^n \frac{u_i(a)}{s(u_i)}.
\end{aligned}$$

Similarly, we may derive an upper bound, recalling the constant $x = s((1, 0, ..., 0))^{-1}$:

$$\frac{\mathsf{sw}(a, \tilde{U})}{\tilde{n}} - \frac{\mathsf{sw}(a, U)}{n} \leq \sum_{i=1}^n g_i \frac{u_i(a)}{s(u_i)} + n\varepsilon x - \frac{1}{n} \sum_{i=1}^n u_i(a) \leq \sum_{i=1}^n \bar{s}_i \frac{u_i(a)}{s(u_i)} + n\varepsilon x - \frac{1}{n} \sum_{i=1}^n u_i(a) = n\varepsilon \cdot x.$$

Since $\varepsilon > 0$ was arbitrary, and since both of the latter two bounds tend to 0 as $\varepsilon \longrightarrow 0$, we can now choose $\varepsilon > 0$ small enough to fulfill all of the inequalities in the Proposition B.8 for any prescribed threshold $\rho > 0$. This proves the claim. □

Our result shows that, from the perspective of worst-case distortion, using a stakes-based recomposition is equivalent to assuming across the population that every voter has equal stakes.

B.9 Formalisms about the STABLE LOTTERY RULE

We now define the STABLE LOTTERY RULE,following [Ebadian et al., 2022]. Since only the case of a stable lottery of size $\sqrt{m}$ is relevant to us, we shall restrict our definition to this special case. Let $\mathcal{P}_{\sqrt{m}}([m])$ be the set of all subsets (or 'committees') of $[m]$, of size $\sqrt{m}$, and let $\Delta(\mathcal{P}_{\sqrt{m}}([m]))$ be the set all of all distributions on $\mathcal{P}_{\sqrt{m}}([m])$. Given a subset $A \subseteq [m]$ of alternatives, an alternative $a \in [m]$ and a histogram profile $\mathbf{h} \in \Delta(S_m)$, let us denote the fraction of voters who rank a ahead of all of A by

$$Freq_{a>A}(\mathbf{h}) = \sum_{\pi \in S_m} h_\pi \mathbb{I}(a >_\pi A).$$

If $a \in A$, then we set $Freq_{a>A}(\mathbf{h}) = 0$ for all $\mathbf{h}$.

Definition B.9 (Stable lottery)**.** Given a preference histogram $\mathbf{h}$, a *stable lottery* (of size $\sqrt{m}$) is a probability distribution $P(\mathbf{h}) \in \Delta(\mathcal{P}_{\sqrt{m}}([m]))$ (i.e., a random selection of a committee of size $\sqrt{m}$) such that for all $\mathbf{h}$,

$$\max_{a \in [m]} \mathbb{E}_{A \sim P(\mathbf{h})} \left[Freq_{a>A}(\mathbf{h}) \right] < \frac{1}{\sqrt{m}}.$$

It is well-known that a stable lottery always exists, see, e.g. [Ebadian et al., 2022]. Building on this definition, we define the STABLE LOTTERY RULE in terms of histograms.

Definition B.10 (STABLE LOTTERY RULE)**.** Given a histogram $\mathbf{h}$, let $P(\mathbf{h})$ be a stable lottery. With probability $1/2$, sample a committee A of size $\sqrt{m}$ from $P(\mathbf{h})$, and then choose an alternative uniformly at random from A. Else, with the remaining probability $1/2$, simply choose an alternative uniformly at random from $[m]$.

PROOF OF THEOREM 3.12.

THEOREM 3.12 (**UPPER BOUND**). *For $s \in \{sum, max, range\}$,* $\mathrm{dist}^s(\text{STABLE LOTTERY}) \in O(\sqrt{m})$.

First, assume that $s \in \{\mathrm{max}, \mathrm{sum}\}$, and let f = STABLE LOTTERY RULE. Then, by a well-established result from Ebadian *et al* [Ebadian et al., 2022], we know that both for $s = \mathrm{sum}$ and $s = \mathrm{max}$, the worst-case distortion over unit-stakes instances is of the order $O(\sqrt{m})$,

$$\sup_{n \geq 1} \sup_{U \in \mathcal{U}_s} \mathrm{dist}_U(\text{STABLE LOTTERY RULE}) \in O(\sqrt{m}),$$

where we recall the notation $\mathcal{U}_s$ for the set of utility matrices U where each voter has unit stakes, $s(\mathbf{u}_i) = 1$. Our goal is to use Theorem B.5 to conclude that the stakes-proportional procedure also has distortion of the order at most $O(\sqrt{m})$. For this, we need to confirm that the STABLE LOTTERY RULE is rationally approximable in the sense of Definition B.6. Indeed, this is seen as follows. Let $\mathbf{h}$ be an arbitrary preference histogram. In [Ebadian et al., 2022], it is proven not just that a stable lottery always exists for $\mathbf{h}$; indeed, a slightly stronger requirement is validated, namely, that the lottery satisfies

$$\max_{a \in [m]} \mathbb{E}_{A \sim P(\mathbf{h})} \left[Freq_{a>A}(\mathbf{h}) \right] \leq \frac{1}{\sqrt{m}+1}.$$

Now, let $\varepsilon > 0$. Suppose that $\tilde{\mathbf{h}}$ is another histogram profile with rational entries such that

$$\sup_{\pi \in S_m} |h_\pi - \tilde{h}_\pi| \leq \varepsilon.$$

We may also choose $\tilde{\mathbf{h}}$ such that the difference $\left| Freq_{a>A}(\mathbf{h}) - Freq_{a>A}(\tilde{\mathbf{h}}) \right| \leq \varepsilon$ for any a. Choosing ε small enough, $P(\mathbf{h})$ is a permissible stable lottery also for $\tilde{h}$. Using this stable lottery, we have that $f(\mathbf{h}) = f(\tilde{h})$; thus f is rationally approximable; the statement follows for $s \in \{\mathrm{max}, \mathrm{sum}\}$.

It remains to show the claim for $s = \text{range}$. Here, we argue along the same lines as Observation B.2: The worst-case distortion both for $s = \text{range}$ and for $s = \max$ can be realized while only considering utility matrices in which each voter has minimum utility 0. Let this set of utilities be denoted by $\mathcal{V}$. Then,

$$\sup_{U \in \mathbb{R}_{\geq 0}^{n \times m}} \text{dist}_U^{\text{range}}(f) = \sup_{U \in \mathcal{V}} \text{dist}_U^{\text{range}}(f) = \sup_{U \in \mathcal{V}} \text{dist}_U^{\max}(f) = \sup_{U \in \mathbb{R}_{\geq 0}^{n \times m}} \text{dist}_U^{\max}(f).$$

□

B.10 Folklore: all randomized rules have at least m distortion.

Fact B.11. *For all voting rules f, dist$(f) \geq m$.*

Proof. Consider a histogram in which each of the m alternatives occupies a $1/m$ fraction of the first positions and the second to last positions are occupied arbitrarily. There exists some alternative a which will be chosen by the randomized rule with probability at most $1/m$. Let G denote the group in which a is ranked first. In this group, let us assign the ordered utility vector $(1, 0, ..., 0)$. In the remainder of the population G^c, we assign the zero utility vector. Let us denote this utility matrix by U. Then, since f selects a with probability at most $1/m$, denoting the winner of the election by a', we obtain $\mathbb{E}[\text{sw}(a', U)/n] \leq 1/m^2$, while the maximum welfare alternative has average utility $\text{sw}(a, U)/n = 1/m$; thus the distortion of f is at least m. □