Accounting for Stakes in Democratic Decisions

BAILEY FLANIGAN, Carnegie Mellon University
ARIEL D. PROCACCIA, Harvard University
SVEN WANG, Massachusetts Institute of Technology

When a society makes a collective decision on a political issue – e.g., votes to decide which policy to enact – different members of the society are inevitably affected to differing degrees by the decision: some people may benefit to the same extent from all alternatives, while others may gain a lot from certain alternatives and lose a lot from others. We formalize this idea as people having different *stakes* in the decision, where someone with *high stakes* stands to gain or lose a lot, depending on the decision's outcome.

There is a mainstream notion that stakeholders—i.e., people with high stakes—should be given sufficient say in collective decisions; however, social choice offers no framework for building democratic decision processes that account for people's differing stakes. In this paper, we build and implement such a framework. First, we introduce a formal approach to *measuring* a voter's stakes. Then, we formalize what it means for a democratic decision process to "account for" these stakes by defining *stakes-based proportionality*: giving voters decision-making power in proportion to their stakes. Within this framework, we formalize the intuition that giving stakeholders sufficient say is important: we show that accounting for stakes can dramatically increase the quality of decisions made through voting, as measured by the *distortion* – the competitive ratio between the utilitarian social welfare of the winner and the highest-welfare alternative. Motivated by these results, we conclude by exploring how to design decision processes that account for stakes.

Finally, we surprisingly find that accounting for stakes is equivalent, from a distortion perspective, to assuming that voters' underlying utilities are uniformly normalized – one of the most common assumptions in the distortion literature.

1 INTRODUCTION

At the backbone of a democratic society are *collective decisions*: society-level choices over, e.g., policies or candidates, on which a society's members weigh in and are in turn potentially affected by the outcome. We start this paper from the fundamental idea that a given collective decision can affect different members of society to *differing degrees*. It is not hard to come up with salient examples: consider, for instance, the decision of whether to instate a city-wide COVID masking requirement. Immunocompromised people and/or essential workers are likely much more affected by the outcome of this decision than people who work from home and can easily shelter in place. In mainstream language, we often talk about these disparities in terms of *stakes*, where someone with *high* stakes in a decision stands to gain or lose significantly depending on its outcome, while someone with *low* stakes is relatively unaffected by the decision.

It is intuitive not only that people's stakes may differ, but that *accounting for* these differing stakes is important to the quality and legitimacy of decisions' outcomes. This can be seen often in critiques of how processes *fail* to do so: for instance, many critique the fact that the global south, despite being disproportionately affected by climate change, has thus far been granted far less power in deciding global climate policy [17]. Political scientific theories engage with this idea, too, in the form of the *most affected principle,* which roughly argues that those most affected by a collective decision deserve greater representation in the process [2, 11]. The importance of accounting for stakes is even codified – if coarsely – in the design of modern electoral democracies: we restrict who can vote in elections based on residency because, e.g., people who reside in Kansas are not impacted by local elections in California.

Authors' addresses: Bailey Flanigan, bflaniga@andrew.cmu.edu, Carnegie Mellon University; Ariel D. Procaccia, Harvard University, arielpro@seas.harvard.edu; Sven Wang, Massachusetts Institute of Technology.

Despite the omnipresent sense that stakes are an important factor in collective decisions, social choice theory – focused precisely on formally studying collective decision processes – offers no framework for even reasoning about people's stakes, let alone designing processes which account for them. Without such a formal framework, we can neither rigorously confirm the intuition that decisions *do* have poor outcomes when people with high stakes are given insufficient voice, nor assert whether – and under what conditions – accounting for stakes can lead to outcomes better for the common good. In this paper, we aim to close these gaps by building a theory of stakes in social choice. We build this theory within perhaps the most ubiquitous collective decision process – voting; however, as we explore in Section 7, our formalisms extend naturally to a much broader set of collective decision processes.

We work within the standard model of ranking-based voting with underlying utilities: its key components are a set of n voters, a set of m alternatives over which these voters must collectively choose, and an $n \times m$ matrix of voters' latent *utilities* U. These utilities measure the extent to which each voter stands to gain from a given alternative being the decision outcome. Voters "vote" by ranking the alternatives according to the order of their utilities. A winner is chosen based on these rankings by a *voting rule* f – a function mapping a set of n such rankings to a single winning alternative (a *deterministic* rule), or a distribution over alternatives (a *randomized* rule). We will focus mostly on deterministic rules, due to their widespread use and the practical challenges associated with deploying randomized rules. We evaluate the quality of the winning alternative by its *utilitarian social welfare*, i.e., the sum of voters' utilities for it. As is standard in this model, we evaluate the entire decision process – instantiated with a specific voting rule – via the *distortion*: the competitive ratio between the welfare of the highest-welfare alternative and that of the winner.

Within this model, we ask and answer the following questions: **Q1.** *Can we formally prove that voting processes that ignore voters' differing stakes have poor outcomes?* **Q2.** *What is the "right" way to **measure** a voter's stakes?* **Q3.** *What does it mean for a voting process to **account for** stakes?* and **Q4.** *In what senses, if any, is accounting for stakes guaranteed to improve voting outcomes?* We see these questions as precursors to asking what might be the ultimate question: *How can we design deployable collective decision processes that account for stakes?* After building up the machinery necessary to approach this key question, in Section 7 we will suggest several ideas for how to implement such processes.

Before describing our approaches and contributions in Section 1.1, we motivate them through Example 1.1, which gives intuition for why the answer to **Q1** is, resoundingly, yes – failing to account for stakes in voting *can* produce arbitrarily poor outcomes for the population. In addition to answering **Q1**, this example will illustrate and motivate our model and give intuition for our approaches to answering **Q2** - **Q4**.

Example 1.1. We consider an election in which neighborhood residents will decide the placement of a new bus stop in their neighborhood. The city has given them two options at distant locations, a and b. The residents fall into two demographic types: residents of the first type, composing 10% of the population, live in the area directly surrounding location a (but are far from b). They do not have cars and rely on the bus, and thus would benefit significantly from a; we therefore say type 1 voters have the utility vector $(10, 0)$ (specifying utilities in the order a, b). The remaining 90% of voters live throughout the rest of the neighborhood – a region encompassing bus stop b. In contrast to type 1 voters, type 2 voters have cars and essentially never take the bus; however, when they *do*, they are inclined to take it from b due to their geographic and demographic separation from the region around a. Their weak preference for b is reflected in their utility vector $(0, \epsilon)$, where $\epsilon > 0$ is some very small number. Based on these utility vectors, 90% of residents will vote for b over a, and any majority-consistent voting rule must choose b as the winner. b intuitively seems like a poor outcome, because the election has chosen a bus stop that virtually no one will ever use, while the other bus stop would have helped a lot of people. This intuitive suboptimality is captured in the fact that social welfare of b – the winner, is more than $1/\epsilon$ times smaller than that of a, the other available alternative.

From a purely theoretical standpoint, this example essentially recovers the known folklore result that, without further assumptions on the voting model, all deterministic voting rules have unbounded distortion.[1] However, *conceptually*, it paints a much richer picture of why stakes matter. Notice that the specific chosen numbers did not matter much: the key cause of the distortion was that type 1 voters had high stakes relative to type 2 voters, but they did not have the voting power to sway the election. Then, when type 2 voters got their way, they did not stand to benefit much, and little value was generated for the population. Given that we people *are* likely to have disparate stakes in real issues – and sometimes minority groups may have much higher stakes than the majority – this example seems practically pressing: it suggests that such welfare loss could occur in real elections. This is especially concerning because, as illustrated in both Example 1.1 and preceding real-world examples, the highest-stakes voters – who stand to lose the most when this loss occurs – are likely to be already-marginalized, having high stakes precisely because they have less resources to adjust to policies that are suboptimal for them.

1.1 Approach and Contributions

1.1.1 A theoretical framework for stakes in voting. In Section 2.2, we embed a formal model of stakes into the standard voting model. To do so, we first formalize how to *measure* a voter's stakes (**Q2**). Example 1.1 gives the intuition that a voter's stakes are captured in their *utility vector*, and further seems to suggest measuring a voter's stakes as the difference between their utilities; however, this becomes less obvious when $m > 2$.[2] We thus define and study general *stakes functions*, which map a voter's utility vector to a scalar quantifying their stakes in the election. We then define a voting process that *accounts for* stakes (with respect to a given stakes function) as one which represents voters in the electorate *in proportion to their stakes* (**Q3**). One intuitive scenario in which a *stakes-proportional electorate* could arise is one where voters decide to submit a vote with probability proportional to their stakes (where stakes are normalized to be in $[0, 1]$). However, there are a multitude of ways to create stakes-proportionality, as we illustrate in Section 7.

1.1.2 Characterizing the distortion when stakes are accounted for. In Sections 4 to 6, we characterize the distortion of voting rules when electorates are stakes-proportional with respect to arbitrary (and sometimes specific) stakes functions (**Q4**). In Section 4, we study the distortion of *deterministic* voting rules. We first give bounds applying to *all* deterministic rules and stakes functions. We then deduce from these bounds that when stakes are measured as a voter's *maximum* utility, the voting rule Plurality – when applied to a stakes-proportional electorate – has the *optimal* distortion m over all possible deterministic voting rules and stakes functions. Section 5 offers a parallel (but less detailed) analysis for randomized rules, which concludes that the Stable Lottery Rule, as introduced by Ebadian et al. [7], paired with the same stakes function, achieves the (essentially) optimal distortion of $\sqrt{m}$. Finally, in Section 6, we prove practically-motivated robustness results. First, we show that the distortion of *all rules* is robust to *approximate* stakes-proportionality – a likely outcome of any practical mechanism. Then, given that distortion is a worst-case notion, we investigate the welfare impact of stakes-proportionality on *individual instances*. This yields an intuitive sufficient condition for stakes-proportionality to improve the welfare, instance-wise.

[1]In this example, any voting rule must either choose a or b. If it chooses a, the distortion is unbounded as above. If it chooses b from this set of rankings, then let voters of types 1 and 2 have utility vectors $(\epsilon, 0)$ and $(0, 1)$ respectively – the distortion is again unbounded. We make this formal in Appendix B.1.
[2]Consider the utility vectors $(1, 1, 0)$ and $(1, 0, 0)$. Which reflects higher stakes?

1.1.3 Generalizing the unit-stakes assumption. Beyond **Q1 - Q4**, in Section 3 we show how the concept of stakes generalizes one of the main assumptions in the literature permitting bounded distortion: that all voters' utility vectors have the same sum [5, 6, 15] (or, similarly have the same maximum [7]). We observe that in our model, these assumptions can be cast as voters having *unit stakes* with respect to different stakes functions. This conceptual generalization leads to a theoretical one: we show that stakes-proportional representation in our model is *equivalent*, from a distortion perspective, to assuming that voters have unit stakes with respect to the same stakes function. A key implication of this equivalence is that it grants practical motivation to bounds achieved in unit-stakes models: while it seems hard to argue that voters will truly have unit stakes in practice, as we explore in the discussion, there is hope of mechanistically creating the equivalent condition of stakes-proportionality.

1.2 Related Work

To our knowledge, stakes have never been modeled – let alone studied – in the social choice literature.[3] It is perhaps unsurprising, then, that most decision processes studied in social choice do not account for stakes; instead, most aim for some notion of *equality* across people, *regardless* of their stakes. For instance, approval voting, ranking-based voting, and liquid democracy [10] allot each person one vote. Even voting rules that elicit cardinal preferences do not necessarily account for stakes: for instance, consider range voting [13], in which voters express their utilities through scores within a bounded range $[0, 1]$: in Example 1.1, even if voters' reported scores reflected their *true relative utilities* for a and b, type 1 and 2 voters would report scores of $(1, 0)$ and $(0, 1)$, and the election outcome would be the same. The voting system that gets the closest to accounting for stakes is quadratic voting (QV) [14], though it does so under several additional assumptions, including that votes are purchased with externally-valuable currency. While the QV literature does not formally engage with the concept of stakes, QV does deviate from the one-person, one-vote model in a stakes-related way: agents will not purchase votes if they do not care about an issue, possibly leading to something resembling stakes-dependent representation, though this requires formal treatment.

Even outside of strictly voting-based mechanisms, we still see the principle of equality: for example, a key desideratum of deliberative processes like citizens' assemblies is to give everyone in the constituency *equal* chance of participation [8]. While many such deliberative processes also aim for proportional demographic representation, this is distinct from – and can even be in tension with – ensuring that stakeholders have sufficient say.[4]

Our work does fall within a category of *extensions* of the voting model, in which the voting rule has some limited additional auxiliary information about the utilities; however, even existing such extensions do not engage with the concept of stakes. Existing work often assumes that this information comes from queries to the utilities (see [4] for an overview); one somewhat stakes-related exception assumes that voters report, per pair of alternatives, whether their ratio of utilities exceeds a threshold [1]. These types of auxiliary information differ fundamentally from our notion of stakes in that, while other types of information must be reported by voters, stakes – capturing in a single number how much a voter cares about the outcome of an election – may be revealed in their behavior (see Section 7 for examples).

[3] We know of two papers in voting that prominently feature the term "stakes" [3, 12]; both use the term differently and explore unrelated questions.

[4] This tension is illustrated by a deliberative poll in Australia, whose goal was to decide how to facilitate reconciliation between indigenous and non-indigenous groups. Indigenous people were a very small fraction of the overall population; thus, giving all constituents an equal chance at a seat at the table would necessitate making indigenous people a very small part of the deliberation body. At the same time, indigenous people were affected to an outsized degree by the decision at hand. In their experiment, Jimenez [11] investigated the impacts of intentionally over-representing indigenous people in some groups of deliberators (i.e., prioritizing stakes at some cost to equality).

2 MODEL

The model proceeds two parts: Section 2.1 establishes the standard voting model with underlying utilities from the distortion literature. Section 2.2 then introduces a framework for studying stakes into the standard definitions. Throughout the paper, we will often talk about vectors containing a string of ℓ ones followed by a string of ℓ' zeros, which we denote by $\mathbf{1}_\ell \mathbf{0}_{\ell'}$. We use $\mathbb{I}(\cdot)$ to denote the indicator function.

2.1 Voting and Distortion

There are n voters and m alternatives. We let $[n]$ the set of all voters, sometimes called the *electorate*, and let $[m]$ be the set of all alternatives. We assume that both voters and alternatives have some fixed numbering, and we refer to individual voters as i and alternatives as a. Each voter i has some nonnegative, real *utility* for each alternative, and we summarize voters' utilities in the *utility matrix* $U \in \mathbb{R}_{\geq 0}^{n\times m}$. We refer to i's vector of utilities over the alternatives as $\mathbf{u}_i \in \mathbb{R}_{\geq 0}^{m}$, corresponding to the the i-th row of U. We let $\mathbf{u}$ refer to an arbitrary utility vector that is not associated with a particular voter. We refer to i's utility for a specific alternative a as $u_i(a) \in \mathbb{R}_{\geq 0}$, corresponding to the a-th entry of $\mathbf{u}_i$. Likewise, $u(a)$ is the utility for a specified by utility vector $\mathbf{u}$. In our problem, U constitutes an *instance.*

2.1.1 Voters' expressed preferences. Each voter i expresses their preferences via a complete *ranking* over (i.e., permutation of) the alternatives, denoted by $\pi_i \in S_m$, where S_m is the set of all permutations of $[m]$. We say i prefers alternative a to a' if a precedes a' in the permutation π_i, as denoted by $a \succ_{\pi_i} a'$. Abusing notation slightly, we will use $\pi(j)$ to denote the alternative ranked in the j-th position in ranking π.

As others have before (e.g., [18]), we summarize a collection of rankings with a preference *histogram*, expressed as the $m!$-length vector $\mathbf{h} = (h_\pi : \pi \in S_m)$. Its π-th entry, $h_\pi \in [0,1]$, is the fraction of rankings in the collection equal to π. As such, $\sum_{\pi\in S_m} h_\pi = 1$, and thus the space of all possible preference histograms is just the space over all valid distributions over S_m. We refer to the simplex of all possible histograms as $\Delta(S_m) := \left\{\mathbf{h} \in [0,1]^{S_m} : \sum_{\pi\in S_m} h_\pi = 1\right\}$. Occasionally, we will instead reason about voters' preferences as a discrete set of n rankings that is *consistent* with a given histogram $\mathbf{h}$ – that is, a set of rankings in which each $\pi \in S_m$ appears $n \cdot h_\pi$ times. We refer to any set of n rankings as a preference *profile* $\boldsymbol{\pi}$, and we define $\Pi^{\mathbf{h}}$ as the set of profiles consistent with a histogram $\mathbf{h}$. Note that if $\mathbf{h}$ contains any irrational entries, then $\Pi^{\mathbf{h}} = \emptyset$; otherwise, $\Pi^{\mathbf{h}}$ contains countably infinite profiles.

2.1.2 Translating utilities into rankings. Each voter's ranking is determined simply by the ordering of their utilities. That is, agent i will order alternatives in decreasing order of their utilities, so that $u_i(a) > u_i(a') \implies a \succ_{\pi_i} a'$ for all $a, a' \in [m]$. [5] A voter's utility vector implies a ranking, and so it follows that an entire utility matrix implies a preference histogram. We denote the histogram implied by a utility matrix U as $\mathrm{hist}(U)$, whose π-th entry is given by

$$\mathrm{hist}_\pi(U) = \frac{1}{n}\sum_{i\in[n]} \mathbb{I}\{\pi_i = \pi\}, \quad \text{for all } \pi \in S_m.$$

2.1.3 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.[6] We will often refer to this class of functions as *randomized* rules to explicitly distinguish them from their sub-class,

[5] In our analyses, we assume worst-case rankings in the case that $u_i(a) = u_i(a')$. This is for simplicity of our lower bounds; one could instead tie-break explicitly in instances by perturbing the utilities by arbitrarily small amounts.

[6] Histograms are inherently anonymized; as such, we will study exclusively anonymous voting rules (encompassing essentially all voting rules).

deterministic voting rules, which map a histogram to a *single* alternative – i.e., for which, for any $\mathbf{h}$, the support of $f(\mathbf{h})$ is of size 1.

Among randomized rules, we consider the Stable Lottery Rule [7], which at a high level draws a winner either at random or from a *stable lottery* – a randomization over a subset of $[m]$ that is, in some sense, preferred by voters to other such subsets. We will not engage directly with this rule's precise definition, so we defer it, for brevity, to Appendix C.2.

Among deterministic rules, we study two main sub-classes: *positional scoring rules* and *Condorcet-consistent* rules. A positional scoring rule is defined by a scoring vector $\mathbf{w} \in [0,1]^m$. Alternative a gets w_j points each time it is ranked jth, so its total score is given by $\sum_{j=1}^{m} \sum_{\pi \in S_m} w_j \cdot h_\pi \cdot \mathbb{I}(\pi(j) = a)$. The winning alternative is the one with the highest score. Our results most strongly feature the positional scoring rule Plurality, defined by the scoring vector $\mathbf{w} = \mathbf{1}_1 \mathbf{0}_{m-1}$. We additionally mention Borda Count, defined by the linearly-decreasing scoring vector $\mathbf{w} = (1, (m-2)/(m-1), \ldots, 1/(m-1), 0)$, and Veto, defined by $\mathbf{w} = \mathbf{1}_{m-1}\mathbf{0}_1$. To define Condorcet-consistent rules, we first establish 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.

2.1.4 Distortion. An instance of our problem consists of a utility matrix U. For a given U, every alternative has some *utilitarian social welfare*, equal to the sum of voters' utilities for that alternative:

$$\mathrm{sw}(a, U) := \sum_{i \in [n]} u_i(a).$$

We will often use a^* to denote the highest social welfare alternative in U, i.e., $a^* := \arg\max_{a \in [m]} \mathrm{sw}(a, U)$.

In the literature, the extent to which a voting rule selects a winner with high social welfare is measured by the competitive ratio between the social welfare of the winner and that of the highest-welfare alternative a^*. For a given rule f, the *instance-specific distortion* $\mathrm{dist}_U(f)$ is equal to this ratio in a specific instance U. As is standard, we evaluate the overall social welfare of f via its *distortion*, $\mathrm{dist}(f)$, which refers to the worst case of this competitive ratio over all possible utility matrices U. Formally, these two notions are defined as follows:

$$\mathrm{dist}_U(f) := \frac{\mathrm{sw}(a^*, U)}{\mathbb{E}[\mathrm{sw}(f(\mathrm{hist}(U)), U)]} \qquad \text{and} \qquad \mathrm{dist}(f) := \sup_{n \geq 1} \sup_{U \in \mathbb{R}_{\geq 0}^{n \times m}} \mathrm{dist}_U(f),$$

where the expectation in $\mathrm{dist}_U(f)$ is taken over the distribution $f(\mathrm{hist}(U))$. In the definition of $\mathrm{dist}(f)$, we take the supremum over n to more conveniently deal with the fact that in worst-case instances, n must be large enough relative to m in order to realize utility matrices with m-dependent fractional compositions. The fact that we do not also take the supremum over m reflects that we consider the distortion to be a *function* of m, as is standard in the distortion literature.

2.2 A stakes framework within the voting model

2.2.1 "Measuring stakes" via stakes functions. A *stakes function* is any map $s : \mathbb{R}_{\geq 0}^m \to \mathbb{R}$ which maps utility vectors to a scalar measure of the stakes associated to that vector. Conceptually, voter's stakes should depend on the relative magnitudes of their utilities for alternatives, but not *which* alternatives they prefer; we thus restrict to functions s which are *permutation invariant*. For example, if $m = 2$, we want voters with utility vectors $(0, 1)$ and $(1, 0)$ to have the same stakes. In particular, the values of s are uniquely determined by the stakes s assigns to utility vectors which are *ordered* in decreasing order; we will use this property throughout the paper.

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, as 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:

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

2.2.2 "Accounting for stakes" via stakes-proportionality. At a high level, 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 duplicating voters in proportion to some function of $s(\mathbf{u}_i)$. While these weights could be defined by any such "recomposition function", we focus here on perhaps the simplest: the unit function, which results in voters being represented *in proportion* to their stakes. Formally, given a stakes function s and a utility matrix U, we let $\mathsf{hist}^s(U)$ be the s-*proportional* histogram arising from U, whose π-th entry is given by

$$\mathsf{hist}^s_\pi(U) = \frac{\sum_{i\in[n]} s(u_i) \cdot \mathbb{I}(\pi_i = \pi)}{\sum_{i\in[n]} s(u_i)}, \qquad \pi \in S_m.$$

In other words, each voter i's ranking is represented in the election with weight $\frac{s(\mathbf{u}_i)}{\sum_{i\in[n]} s(\mathbf{u}_i)}$. When we are not discussing a specific histogram with this property, we will sometimes instead refer to the *stakes-proportional electorate*. In Section 7, we discuss going beyond stakes proportionality to study more general stakes-dependent recomposition functions.

2.2.3 Distortion under stakes-proportionality. By its standard definition, the distortion measures the extent to which the outcome of a voting rule f can be distorted, from a welfare perspective, due to f seeing only the frequency of ordinal rankings in the histogram $\mathsf{hist}(U)$, rather than the underlying cardinal utilities in U. We will now define an analogous competitive ratio, called the s-*distortion*, for when f instead sees the s-proportional histogram $\mathsf{hist}^s(U)$. As with the standard distortion, we define both instance-specific and worst-case notions. For any U and s, let

$$\mathsf{dist}^s_U(f) = \frac{\max_{a\in[m]} \mathsf{sw}(a, U)}{\mathbb{E}[\mathsf{sw}(f(\mathsf{hist}^s(U)), U)]}, \qquad \text{and} \qquad \mathsf{dist}^s(f) = \sup_{n\geq 1} \sup_{U\in\mathbb{R}^{n\times m}_{\geq 0}} \mathsf{dist}^s_U(f).$$

3 STAKES-PROPORTIONALITY $\iff$ THE UNIT-STAKES ASSUMPTION

In the distortion literature, a popular assumption is that voters' utilities are normalized to sum to 1 (or, in recent work, have maximum utility 1). We observe, first, that these assumptions amount to assuming voters have *unit stakes* with respect to the stakes function sum (or respectively, max). This observation illuminates the vast space of possible such s-*unit-stakes* assumptions, where each assumes utilities are normalized according to a different stakes function s. This inspires the open question: *would assuming unit stakes with respect to a different stakes function permit better distortion bounds?* We now close this question for *all voting rules* and *all 1-homogeneous stakes functions* via a surprising result: the s-distortion is *equal to* the distortion when U is restricted to satisfy the s-unit stakes assumption. Combined with results proven later in the paper, this bi-directional reduction will allow us to conclude that among s-unit stakes assumptions, it is optimal to assume max-unit stakes, and that this assumption permits at best distortion m for the rule Plurality. This result also allows us to pass upper and lower bounds freely across the two models, which we will make use of later in the paper.

At a high level, we prove the reductions between the two settings via the constructions illustrated in Figure 1, which shows how to transform an instance of one setting into the other while preserving the distortion. The transformation *from* s-unit-stakes instances is direct: since all voters' stakes are already the same, accounting for stakes has no effect and the distortion is immediately preserved. To produce an instance with s-unit stakes from one without, we simply rescale voters' utility vectors and duplicate them according to their stakes, again preserving the distortion. This rescaling is enabled by the 1-homogeneity of our stakes function s.

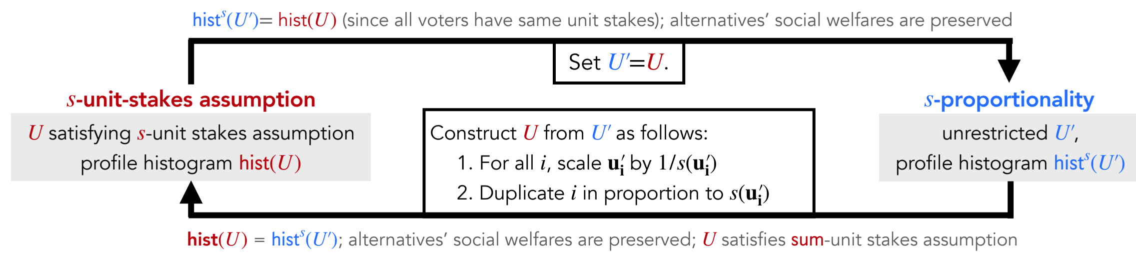

Fig. 1. Constructions giving reductions between the s-unit stakes assumption (existing model) and s-proportionality (our model).

Although the intuition behind this reduction is quite simple, some technicalities arise when $\text{hist}^s(U)$ can have irrational entries. Thus, we first state and prove the result assuming $\text{hist}^s(U)$ has only rational entries, which we ensure by restricting to rational utility matrices $U \in \mathbb{Q}_{\geq 0}^{n\times m}$ and *rationality-preserving* stakes functions s (i.e., $s(\mathbf{u}) \in \mathbb{Q}$ whenever $\mathbf{u} \in \mathbb{Q}_{\geq 0}^{m}$). The proof of Theorem 3.1 is found in Appendix A.1.

Theorem 3.1. *Let f be any voting rule, let s be 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,*

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

While these are already very mild restrictions on U and s, we prove in Appendix A.2 (at the cost of very mild technical conditions on f) the analogous results for real-valued utilities and arbitrary 1-homogeneous stakes functions.

4 DISTORTION OF DETERMINISTIC VOTING RULES

In this section, we analyze the s-distortion of deterministic voting rules. To contextualize our bounds in contrast to what is possible without accounting for stakes, we recall the folklore result that for all deterministic voting rules f, $\text{dist}(f) = \infty$. For completeness we prove this in Appendix B.1.

4.1 Upper Bounds

We now upper-bound the s-distortion of f, for arbitrary s and f. We reason about all voting rules at once by expressing our bounds in terms of the parameter β_f, the minimum fraction of voters that must rank the winner by f in the first position. Formally, if $f(\mathbf{h})$ is the winner of the election summarized by histogram $\mathbf{h}$, then β_f is given by

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

We reason about all stakes functions at once via two key coefficients, $\kappa^{\text{upper}}(s)$ and $\kappa^{\text{lower}}(s)$, which are defined for stakes function s as follows. Intuitively, $\kappa^{\text{upper}}(s)$ and $\kappa^{\text{lower}}(s)$ measure the extent to which a stakes function s can over- or under-estimate a voter's maximum utility.

$$\kappa^{\text{upper}}(s) := \sup_{\mathbf{u}\in\mathbb{R}^m_{\geq 0}} \frac{s(\mathbf{u})}{\max(\mathbf{u})} \quad \text{and} \quad \kappa^{\text{lower}}(s) := \inf_{\mathbf{u}\in\mathbb{R}^m_{\geq 0}} \frac{s(\mathbf{u})}{\max(\mathbf{u})}. \tag{1}$$

We now upper-bound the s-distortion using the key insight that β_f and the κ coefficients are linked: β_f captures how frequently the winner is ranked first, and the κ's capture how closely s captures a voter's highest utility – which is precisely their utility for their first-ranked alternative. Via connection, we can lower-bound social welfare of the winner.

Theorem 4.1. *For all voting rules f and all stakes functions s,*

$$\mathit{dist}^{\mathbf{s}}(f) \leq \frac{1}{\beta_f} \cdot \frac{\kappa^{\text{upper}}(s)}{\kappa^{\text{lower}}(s)}$$

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

$$\text{sw}(a,U) = \sum_{i\in[n]} u_i(a) \leq \sum_{i\in[n]} \max_a u_i(a) \leq \frac{\sum_{i\in[n]} s(\mathbf{u}_i)}{\kappa^{\text{lower}}(s)} \qquad \text{for all } a \in [m]. \tag{2}$$

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':

$$u_i(a') = \max_{a\in[m]} u_i(a) \geq \frac{s(\mathbf{u}_i)}{\kappa^{\text{upper}}(s)}. \tag{3}$$

Also, $N_{a'}$ must compose at least a β_f fraction of the stakes-proportional electorate, else a' would not be the winner:

$$\frac{\sum_{i\in N_{a'}} s(\mathbf{u}_i)}{\sum_{i\in[n]} s(\mathbf{u}_i)} \geq \beta_f.$$

This fact, combined with equation (3), gives a lower bound on a''s social welfare:

$$\text{sw}(a',U) \geq \sum_{i\in N_{a'}} u_i(a') \geq \frac{\sum_{i\in N_{a'}} s(\mathbf{u}_i)}{\kappa^{\text{upper}}(s)} \geq \frac{\beta_f \sum_{i\in N} s(\mathbf{u}_i)}{\kappa^{\text{upper}}(s)}.$$

Combining this with equation (2) and denoting the maximum welfare alternative by a^*, we obtain that

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

We remark that in the definitions of $\kappa^{\text{upper}}(s)$ and $\kappa^{\text{lower}}(s)$, we could have replaced max with range and the upper bound would still hold. This is because the worst-case distortion can always be realized over only utility matrices in which each voter's minimum utility is 0, in which case max = range. We formalize this argument in Appendix B.2.

4.2 Lower Bounds

Now, we ask: *what is the minimum possible s-distortion achievable by any deterministic voting rule, for any stakes function s?* Theorem 4.2, our main lower bound, shows that no deterministic rule can achieve better distortion than $m-1$.

Theorem 4.2. *For all deterministic voting rules f and all stakes functions s, $\mathit{dist}^{\mathbf{s}}(f) \geq m-1$.*

The proof is found in Appendix B.3, and proceeds by constructing two instances U, U', both in which all voters have identical utility vectors and thus identical stakes for all s (so that $\text{hist}(U) = \text{hist}^s(U)$ for all s, and likewise for U'). Then, we show that all voting rules must have distortion at least $m-1$ in at least one of the instances.

One may notice that when $\beta_f = 0$, there is an unbounded gap between the lower bound in Theorem 4.2 and the upper bound in Theorem 4.1. We now find that in this case, unfortunately the distortion is indeed unbounded:

Proposition 4.3. *Let f be any deterministic voting rule with $\beta_f = 0$ and any stakes function s, $dist^s(f) = \infty$.*

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, $\mathsf{dist}_U(f)$ is unbounded. Because all voters have identical utility vectors, $\mathsf{hist}(U) = \mathsf{hist}^s(U)$, implying $f(\mathsf{hist}(U)) = f(\mathsf{hist}^s(U))$; hence $\mathsf{dist}^s_U(f)$ is also unbounded. □

This lower bound is practically significant because most popular voting rules have $\beta_f = 0$: a simple instance shows this to be true for all Condorcet-consistent voting rules, as well as Borda and Veto.[7] This negative result motivates our study of Plurality, whose β_f *is* lower-bounded, and as we will show, *maximal* over all deterministic voting rules.

4.3 Plurality and its Optimality

To prove this optimality, we apply Theorem 4.1 to upper-bound $\mathsf{dist}^{\max}(\text{Plurality})$ by m. We minimize this upper bound by choosing the voting rule f which maximizes β_f, and then separately choosing the stakes function s which minimizes $\kappa^{\text{upper}}(s)/\kappa^{\text{lower}}(s)$. We first show that Plurality maximizes β_f amongst all deterministic rules:

Lemma 4.4. *For any voting rule f, $\beta_f \leq 1/m$. Moreover, $\beta_{\text{Plurality}} = 1/m$.*

Proof. The fact that $\beta_f \leq 1/m$ for any f is proven in Lemma B.3. The fact that $\beta_{\text{Plurality}} \geq 1/m$ follows immediately from the definition of Plurality: 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. □

Now, to select the s that minimize $\kappa^{\text{upper}}(s)/\kappa^{\text{lower}}(s)$, we first observe that for all s, $\kappa^{\text{upper}}(s)/\kappa^{\text{lower}}(s) \geq 1$. Therefore, setting $s = \max$, which satisfies $\kappa^{\text{upper}}(\max)/\kappa^{\text{lower}}(\max) = 1$, attains the minimal possible value for this ratio. Theorem 4.1 then immediately implies the following upper bound on the max-distortion of Plurality (and likewise for the range-distortion, by the reasoning at the end Section 4.1). We conclude the desired corollary:

Corollary 4.5. *$dist^{max}(\text{Plurality}) \leq m$ and $dist^{range}(\text{Plurality}) \leq m$.*

Since these upper bounds match our lower bound in Theorem 4.2, we conclude that Plurality, under max- or range-proportionality, achieves *optimal* s-distortion over *all* deterministic voting rules and *all possible* stakes functions.

4.3.1 A finer-grained lower bound for Plurality. 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. As a result, one may want to understand the s-distortion of Plurality for *general* s. To this end, we provide Theorem 4.6, a lower bound on the s-distortion of Plurality for any s. This bound essentially matches the upper bound in Theorem 4.1, except it is in terms of a slightly modified version of κ^{lower}, called $\tilde{\kappa}^{\text{lower}}$. $\tilde{\kappa}^{\text{lower}}$ is again the infimum of the ratio $s(\mathbf{u})/\max(\mathbf{u})$, except now over only utility vectors whose largest two entries are identical. For most s, the gap between $\kappa^{\text{lower}}(s)$ and $\tilde{\kappa}^{\text{lower}}(s)$ is either nonexistent, or at most a factor of 2. We formally define $\tilde{\kappa}^{\text{lower}}$ and prove the theorem in Appendix B.5. In the

[7] Indeed, consider the following instance with 4 alternatives, a, b, c, d. Split the voters into three equal-sized groups, with the three groups having rankings $b \succ a \succ c \succ d$, $c \succ a \succ b \succ d$, and $d \succ a \succ b \succ c$, respectively. Then, a is ranked ahead of any other alternative in 2/3 of voters, and is the Condorcet winner; it will also be the Borda winner. However, it is never ranked first.

construction, we let the highest-welfare alternative be ranked second by many voters, where – due to our use of $\tilde{\kappa}^{\text{lower}}$ instead of κ^{lower} – it can amass substantial utility while not gaining sufficient PLURALITY points to win.

THEOREM 4.6. *For any stakes function s and proportional recomposition, we have that*

$$dist^s(\text{PLURALITY}) \geq (m-1)\frac{\kappa^{\text{upper}}(s)}{\tilde{\kappa}^{\text{lower}}(s)}.$$

Since sum-unit stakes is commonly assumed by the literature – and, per Section 3, any bounds attained in our model translate to theirs – we now illustrate how we can easily apply our bounds on the distortion of PLURALITY to recover existing results. Via a straightforward application[8] of Theorems 4.1 and 4.6, we find that $m^2/2 \leq \text{dist}^s(\text{PLURALITY}) \leq m^2$, thereby recovering the existing result that, assuming sum-unit stakes, PLURALITY has $\Theta(m^2)$ distortion (see Theorem 1 of [5]). In fact, our bounds are tighter, improving upon the gap in their bounds from a factor of 8 to a factor of 2.

5 DISTORTION OF RANDOMIZED VOTING RULES

Across the distortion literature, randomized rules can achieve lower distortion than deterministic rules, and a natural question is whether – and to what extent – this phenomenon persists under stakes-proportionality. To explore this question, we first prove a lower bound showing that, for all 1-homogeneous stakes functions s, any voting rule must suffer s-distortion at least $\Omega(\sqrt{m}/\log m)$. The construction is fairly intricate, so we defer it to Appendix C.1.

THEOREM 5.1. *For all randomized voting rules f and all 1-homogeneous stakes functions s,*

$$dist^s(f) \geq \frac{\sqrt{m}}{10 + 3\log m}.$$

As we did for deterministic rules, we now aim to identify a stakes function-voting rule pair whose s-distortion matches this lower bound. In order to do so, we use our reduction in Section 3 to carry over an existing result from the unit stakes model. In particular, Ebadian et al. [7] show that the STABLE LOTTERY RULE (Definition C.2) has distortion at most $\sqrt{m}$ assuming either sum-unit stakes or max-unit stakes. Per their Theorem 3.4 combined with our reduction, we conclude Corollary 5.2, which extends also to range by the same intuition as in Section 4.1. We give this corollary slightly more formal treatment in Appendix C.2.

Corollary 5.2. *Let $s \in \{sum, max, range\}$. Then, it holds that $dist^s(\text{STABLE LOTTERY RULE}) \in O(\sqrt{m})$.*

From these findings, we take away two things: first, randomized rules can indeed dominate deterministic rules in terms of s-distortion, and by at least a factor of $\sqrt{m}$ (Corollary 5.2 and Theorem 4.2). Second, accounting for stakes can decrease the distortion of randomized rules by at least a factor $\sqrt{m}$, as implied by Corollary 5.2 along with the known fact that *without* accounting for stakes, all randomized voting rules suffer at least distortion m (for completeness, we prove this in Appendix C.3).

6 ROBUSTNESS OF DISTORTION BOUNDS

Here, we consider two practically-motivated forms of robustness. First, we consider robustness to *approximate* stakes-proportionality. Second, we ask whether – and under what conditions on instances – stakes-proportionality not only controls the worst-case distortion, but decreases the welfare loss on a *per-instance basis*.

[8]To apply these bounds, we use that $\beta_{\text{PLURALITY}} = 1/m$, $\kappa^{\text{upper}}(\text{sum}) = m$ per $\mathbf{u} = \mathbf{1}$, $\kappa^{\text{lower}}(\text{sum}) = 1$ per $\mathbf{u} = \mathbf{1}_1\mathbf{0}_{m-1}$, and $\tilde{\kappa}^{\text{lower}}(\text{sum}) = 2$.

6.1 Distortion under approximate stakes-proportionality

Now, suppose we achieve s-proportionality according to *slightly incorrect* values of each voter i's stakes $\hat{s}(\mathbf{u}_i)$, where this estimate is bounded within some $\delta \geq 1$ factor of i's true stakes – that is, $\hat{s}(\mathbf{u}_i) \in [s(\mathbf{u}_i), \delta_i \cdot s(\mathbf{u}_i)]$ for some $\delta_i \in [1, \delta]$.[9] We index δ_i per voter to indicate that this error can differ adversarially across voters, and let $\boldsymbol{\delta} \in [1, \delta]^n$ be a vector of these 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

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

The δ, s-distortion of f is then given by

$$\mathsf{dist}^{\delta,s}(f) = \sup_{n \geq 1,\ U \in \mathbb{R}^{n \times m}_{\geq 0},\ \boldsymbol{\delta} \in [1,\delta]^n} \frac{\max_a \mathsf{sw}(a, U)}{\mathbb{E}[\mathsf{sw}(f(\mathsf{hist}^{\boldsymbol{\delta},s}(U), U)]}.$$

Theorem 6.1 proves strong robustness to this type of errors: for all 1-homogeneous stakes functions s, the δ, s-distortion of *any voting rule* exceeds the s-distortion by at most a factor of δ. The formal proof of the theorem is in Appendix D.1, but the intuition is simple: because s is 1-homogeneous, mis-estimating i's stakes by at most δ is the same as overestimating voters' *utilities* by at most δ. Such mis-estimations can change the distortion by at most a factor of δ.

Theorem 6.1. *Let f be any voting rule and let s be any 1-homogeneous stakes function. Then, for any $\delta \geq 1$,*

$$\mathit{dist}^{\delta,s}(f) \leq \delta \cdot \mathit{dist}^{s}(f).$$

6.2 The per-instance impact of stakes-proportionality

Our results so far show conclusively that stakes-proportionality can be powerful in decreasing the distortion; this is particularly true for Plurality, for which max-proportionality brings the distortion from unbounded to m. However, distortion is a worst-case notion; when deciding whether to deploy mechanisms that facilitate stakes-based proportionality, it is also important to know how doing so will affect the welfare *in a given instance*. Our lower bounds in Theorems 4.2 and 5.1 leave open an enticing possibility: that for some s and f, s-proportionality will decrease the distortion in *all instances* relative to the standard electorate – that is, $\mathsf{dist}^s_U(f) \leq \mathsf{dist}_U(f)$ for all U. However, Theorem 6.2 proves that this ideal case is impossible for Plurality and all 1-homogeneous s. Moreover, as we will illustrate below, the argument that gives this bound seems generalizable to other rules, hinting that in general, s-proportionality decreasing the distortion instance-wise may be too much to hope for.

Theorem 6.2. *Let s be a 1-homogeneous stakes function such that $s(\mathbf{1}_2, \mathbf{0}_{m-2})$ is positive. Then, there exists a U such that $\mathit{dist}^s_U(\text{Plurality}) \geq (m-1) \cdot \mathit{dist}_U(\text{Plurality})$.*

We prove this theorem in Appendix D.2. The argument translates a worst-case distortion lower bound into an instance-wise lower bound via the following procedure: we let an $1 - \alpha$ fraction of voters rank a^* first so that it wins the standard election, and the distortion is 1. Then, among the remaining α fraction of voters, we realize the original lower-bound instance, downscaled. We scale the utilities so that voters in the former group have arbitrarily small stakes compared to the latter group; then, under s-proportionality, the former group disappears and the worst-case lower bound instance is recovered. This transformation for turning worst-case lower bounds into instance lower bounds is not specific to Plurality, suggesting that such a result could be proven for a much larger class of voting rules through

[9] Note that we can realize any type of errors with $\delta \geq 1$, because the composition of the resulting electorate is relative.

the transformation of Theorem 4.2 and/or Theorem 5.1. We leave this to future work, because such a general argument requires intricate reasoning to ensure that for all f, our transformation preserves the winner.

6.2.1 Sufficient conditions. Because stakes-proportionality can sometimes increase the distortion, it would be nice to have a sufficient condition on U under which stakes-proportionality is *guaranteed* to decrease the distortion. Here, we give such a condition for the case when $f = \text{Plurality}$ and $s = \max$: that U contains an *affected minority*, where this minority (1) forms a small enough fraction of voters to have limited voting power; (2) has high stakes relative to the rest of voters; and (3) benefits from the alternatives preferred by the majority at most as much as the majority does.

We formally state and prove the sufficiency of this condition in Appendix D.3. The intuition recalls Example 1.1, but this condition is substantially more general because it permits arbitrary m and allows voters within (and outside) the affected minority to have diverse preferences. This condition is also *non-trivial*, in that it implies neither that (a) the standard electorate elects the worst alternative nor (b) the stakes-proportional electorate elects the best alternative. Finally, we remark that this condition is at least roughly observable: in fact, a conceptually similar condition justified the over-representation of indigenous people in the Australian assembly mentioned in the related work [11].

We see two future directions in which additional sufficient conditions can paint a richer picture of the impact of accounting for stakes. First, there may be sufficient conditions for voting rules other than Plurality, offering a positive direction among stark impossibilities when considering the worst-case s-distortion. We also conjecture that there may exist stronger sufficient conditions where the standard electorate violates *stability conditions*, akin to the classic game theoretic notion of the *core* [16]. At a high level, a stability condition would be violated when a subgroup of voters would prefer to take their allotted voting power – in our setting, proportional to their stakes – and run their own election.

7 DISCUSSION

We see our work in this paper as just the beginning of a stakes-based theory of social choice. To lay the groundwork for this area of research, we now provide an extensive discussion of several extensions and threads of potential future work.

7.1 Mechanism design for accounting for stakes

Direct mechanisms for stakes-proportionality. In some contexts, we may have a verifiable *proxy* for estimating stakes – for instance, in Example 1.1, such a proxy might be how often someone takes the bus. If we have such a proxy, we can account for stakes directly by reweighting voters' rankings. Per our robustness results in Section 6, we can do this with certainty that the distortion will be robust to errors in our proxy. In some contexts, reweighting votes could be too politically unpopular; in such cases, a slightly more distant approach would be to establish a decision-making body composed of a subset of voters, and impose stakes-based representation requirements on this assembly.

Issue-tradeoff mechanisms for stakes-proportionality. In the realistic case where we do not have a proxy for stakes, we can rely on the fact that while stakes are a theoretical property of unobservable utilities, they capture how much a voter *cares* about the outcome of an election, and thus may translate to observable behavior. The generative model discussed briefly in Section 1.1.2, in which voters decide whether to vote with probability proportional to their stakes, is an example of how stakes may influence voters' behavior; however, it is likely too optimistic to be relied upon. Instead, we can take a mechanism design approach that makes voters trade off voting power *between elections*, for example by giving them a budget of votes B to allocate across k elections. Intuitively, voters will allocate their votes to the elections in which they have the highest stakes; to formalize this idea and understand when it leads to stakes-proportionality, we must characterize, in a model of utility-optimizing voters, what is a "rational" stakes function. A key challenge here lies in setting these budgets: if they are uniform across voters, then high distortion can still occur when voters have significantly different *total stakes across issues,* motivating the design of ballots in which this is not the case.

Beyond proportional recomposition. As mentioned in Section 2.2, there are a multitude of stakes-dependent ways to recompose an electorate beyond proportionally, as we have here. We can easily extend our model to encapsulate such more general recompositions; let $r : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ be a generic recomposition function; then, we can define the (r, s)-recomposed histogram arising from U such that its π-th entry is

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

Because both our main lower bounds (Theorems 4.2 and 5.1) use instances in which all voters have identical stakes, these lower bounds extend to *any recomposition in the above class*, implying that more sophisticated recompositions cannot be used to go beyond the lower bounds we have established. However, recompositions with, e.g., submodular stakes-dependence might be practically desirable – or might result from certain mechanisms.

Stakes auditing. Another way to apply the stakes framework could be what we term *stakes-auditing*: evaluating collective decision processes in specific instances *ex-post* to understand whether they sufficiently accounted for stakes. In this way, our framework offers a foundation for *formally* establishing policy critiques like the one mentioned in Section 1 about climate policy. Stakes auditing could involve, for example, statistically testing whether some threshold of stakes-proportionality was reached with high probability.

7.2 Studying stakes in other decision processes

In Section 1.2, we outlined how most collective decision processes studied in social choice expressly do not account for voters' underlying stakes. However, as we will illustrate here, *they can be easily adapted to do so.*

To replicate our analysis of voting for a different decision process, model the process as again consisting of n voters deciding over m alternatives, with underlying utilities U. We immediately inherit *stakes functions* for measuring voters' stakes in the decision, and we can again measure the quality of the outcome based on the distortion. The main thing left to define, then, is *stakes-proportionality.* Its formal specification depends on how the process intakes voters preferences and aggregates them; however, the notion is always conceptually the same, amounting to simply giving people "representation" in proportion to their stakes. For example, in liquid democracy, stakes-proportionality could mean allotting votes to delegate in proportion to voters' stakes; in deliberative democracy, this could mean ensuring the deliberative panel is proportionally-representative with respect to stakes. There is a vast space of research questions across the many collective decision processes, because accounting for stakes could have entirely different implications in each, and the mechanisms that enable accounting for stakes could look entirely different.

Zooming out even further, the concept of stakes is so simple and so ubiquitous that it likely underlies a much broader set of applications beyond collective decision processes. For example, we can think about how people's stakes varying geographically – a plausible condition, given patterns of segregation – can impact the welfare effects of different approaches to redistricting or facility location. Really, stakes can matter in any system making decisions that affect people to varying degrees, and in which there are trade-offs that prevent giving everyone their preferred outcome: potential such settings might include, e.g., prediction algorithms for allocating a limited resource.

REFERENCES

[1] Ben Abramowitz, Elliot Anshelevich, and Wennan Zhu. 2019. Awareness of voter passion greatly improves the distortion of metric social choice. In *Web and Internet Economics: 15th International Conference, WINE 2019, New York, NY, USA, December 10–12, 2019, Proceedings 15*. Springer, 3–16.

[2] Afsoun Afsahi. 2022. Towards a principle of most-deeply affected. *Philosophy & Social Criticism* 48, 1 (2022), 40–61.

[3] Jørgen Juel Andersen, Jon H Fiva, and Gisle James Natvik. 2014. Voting when the stakes are high. *Journal of Public Economics* 110 (2014), 157–166.

[4] Elliot Anshelevich, Aris Filos-Ratsikas, Nisarg Shah, and Alexandros A Voudouris. 2021. Distortion in social choice problems: The first 15 years and beyond. *arXiv preprint arXiv:2103.00911* (2021).

[5] Ioannis Caragiannis, Swaprava Nath, Ariel D Procaccia, and Nisarg Shah. 2017. Subset selection via implicit utilitarian voting. *Journal of Artificial Intelligence Research* 58 (2017), 123–152.

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

[7] Soroush Ebadian, Anson Kahng, Dominik Peters, and Nisarg Shah. 2022. Optimized distortion and proportional fairness in voting. In *Proceedings of the 23rd ACM Conference on Economics and Computation*. 563–600.

[8] Bailey Flanigan, Paul Gölz, Anupam Gupta, Brett Hennig, and Ariel D Procaccia. 2021. Fair algorithms for selecting citizens' assemblies. *Nature* 596, 7873 (2021), 548–552.

[9] Bailey Flanigan, Ariel D. Procaccia, and Sven Wang. 2023. Distortion Under Public-Spirited Voting. (2023).

[10] Paul Gölz, Anson Kahng, Simon Mackenzie, and Ariel D Procaccia. 2021. The fluid mechanics of liquid democracy. *ACM Transactions on Economics and Computation* 9, 4 (2021), 1–39.

[11] Gladys Jimenez. 2009. Australia deliberates: 2001 deliberative poll. (2009).

[12] Anson Kahng and Gregory Kehne. 2022. Worst-Case Voting When the Stakes Are High. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 5100–5107.

[13] Marcus Pivato. 2014. Formal utilitarianism and range voting. *Mathematical Social Sciences* 67 (2014), 50–56.

[14] Eric A Posner and E Glen Weyl. 2015. Voting squared: Quadratic voting in democratic politics. *Vand. L. Rev.* 68 (2015), 441.

[15] A. D. Procaccia and J. S. Rosenschein. 2006. The Distortion of Cardinal Preferences in Voting. In *10th*. 317–331.

[16] Herbert E Scarf. 1967. The core of an N person game. *Econometrica: Journal of the Econometric Society* (1967), 50–69.

[17] Erika Strazzante, Stéphanie Rycken, and Vanessa Winkler. 2022. Global North and Global South: How Climate Change Uncovers Global Inequalities. *Generation Climate Europe* (2022).

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

A OMITTED PROOFS FROM SECTION 3

A.1 Proof of Theorem 3.1

THEOREM 3.1. *Let f be any voting rule, let s be 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,*

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

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, $\mathrm{dist}_U(f) = \mathrm{dist}_U^s(f)$. 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 $U \in \mathbb{Q}_{\geq 0}^{n\times m}$, and construct a unit-stakes utility matrix $\tilde{U}$ such that $\mathrm{dist}_{\tilde{U}}(f) = \mathrm{dist}_U^s(f)$. 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 $\tilde{U} \in \mathbb{Q}_{\geq 0}^{\tilde{n}\times m}$ 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 $G_1, \ldots G_n$.
- 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. □

A.2 Extension of Theorem 3.1 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 3.1 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 A.1 (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 A.2. *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 A.2. 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 $\mathsf{dist}_U(f) = \mathsf{dist}^s_U(f)$. 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 $\mathsf{dist}^s_U(f) = \infty$ and there remains nothing to prove. By Proposition A.3, given *any* $\rho > 0$ we may choose a unit-stakes utility matrix $\tilde{U} \in \mathbb{R}^{\tilde{n} \times m}_{\geq 0}$ 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

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

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

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

which proves the claim. □

Proposition A.3 (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,*

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

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

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

Proof. Let $\varepsilon > 0$ be arbitrary and fix any U. By Definition A.1, 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

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

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 $\pi \in S_m$,

$$\sum_{i \in n} g_i \mathbb{I}(\pi_i = \pi) \leq \tilde{\mathbf{h}}_\pi. \tag{4}$$

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 $G_1, ..., G_n$, none of the rankings $\pi \in S_m$ was overrepresented compared to $\tilde{\mathbf{h}}$ – see equation (4) – 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 $\mathrm{dist}_{\tilde{U}}(f)$ induced by $\tilde{U}$ approximates the distortion $\mathrm{dist}_U^s(f)$ 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

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

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

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

Since $\varepsilon > 0$ was arbitrary, and since both of the latter two bounds tend to 0 as $\varepsilon \to 0$, we can now choose $\varepsilon > 0$ small enough to fulfill all of the inequalities in the Proposition A.3 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 OMITTED PROOFS FROM SECTION 4

B.1 Folklore: all deterministic rules have unbounded distortion

Fact B.1. *For any deterministic voting rule f, dist$(f) = \infty$.*

Proof. Fix f and an arbitrarily small constant $\epsilon > 0$. Let $m = 2$ with alternatives a and b, and let n be an arbitrary even number. Define the ranking profile σ in which half of voters rank $a \succ b$ (type 1) and the other half rank $b \succ a$ (type 2). Construct two possible underlying utility matrices U and $\tilde{U}$:

- Let U such that all voters of type 1, type 2 have utility vectors $\mathbf{u}_i^{(1)} = (2\epsilon, 0)$ and $\mathbf{u}_i^{(2)} = (0, 2)$, respectively.
- Let $\tilde{U}$ such that all voters of type 1, type 2 have utility vectors $\tilde{\mathbf{u}}_i^{(1)} = (2, 0)$ and $\tilde{\mathbf{u}}_i^{(2)} = (0, 2\epsilon)$, respectively.

Observe that $\mathrm{hist}(U) = \mathrm{hist}(U')$. At the same time, in U b has significantly higher social welfare than a ($\mathrm{sw}(a, U) = \epsilon$ and $\mathrm{sw}(b, U) = 1$) whereas in $\tilde{U}$ a has significantly higher social welfare than b ($\mathrm{sw}(a, \tilde{U}) = 1$ and $\mathrm{sw}(b, \tilde{U}) = \epsilon$). Because f must select the same winner across these two instances, it must suffer $1/\epsilon$ distortion in one of the instances. □

B.2 Theorem 4.1 holds when κ's are defined with range instead of max

Here we prove the following observation:

Observation B.2. *The bound in Theorem 4.1 remains true also for a slightly different definition of the coefficients $\kappa^{\mathrm{lower}}(s)$, $\kappa^{\mathrm{upper}}(s)$ where max$(\cdot)$ is replaced by range$(\cdot)$,*

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

Proof. Let $U \in \mathbb{R}^{m\times n}_{\geq 0}$ 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

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

Then, we may restrict the arguments in the proof of Theorem 4.1 to utility vectors with zero minimum entry. This leads to a bound where we may use, instead of $\kappa^{\mathrm{upper}}(s)$ and $\kappa^{\mathrm{lower}}(s)$

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

in place of $\kappa^{\mathrm{upper}}(s)$ and $\kappa^{\mathrm{lower}}(s)$. We may further upper and lower bound these last two quantities, respectively, by

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

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

and we then in particular obtain a distortion upper bound with the two expressions on the right hand side in place of $\kappa^{\mathrm{upper}}(s)$ and $\kappa^{\mathrm{lower}}(s)$. □

B.3 Proof of Theorem 4.2

Theorem 4.2. *For all deterministic voting rules f and all stakes functions s, $dist^{\mathbf{s}}(f) \geq m-1$.*

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 $a_1, \ldots, a_{m-1}$. 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 $\pi_i = a_\ell \succ a' \succ a_1 \succ \cdots \succ a_{m-1}$. We display their 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_1	$\succ \ldots \succ$	a_{m-1}
sorted $\mathbf{u}_i$ for $i \in$ group ℓ:	1		1		0	$\ldots$	0
sorted $\mathbf{u}_i'$ for $i \in$ group ℓ:	1		0		0	$\ldots$	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,

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

□

B.4 Proof that $\beta_f \leq 1/m$ for all deterministic f

Lemma B.3. *For any deterministic voting rule f, it holds that $\beta_f \leq 1/m$.*

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 [9], it is shown that for any voting rule f, we have that

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

where Minimax is the voting rule which chooses the alternative a that suffers the least severe worst pairwise defeat; see [9] 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 claim. □

B.5 Proof of Theorem 4.6

Before proving the theorem, we formally define $\tilde{\kappa}^{\text{lower}}$ as

$$\tilde{\kappa}^{\text{lower}} = \inf_{u \in \mathcal{U}} \frac{s(u)}{\max \mathbf{u}}, \qquad \mathcal{U} := \left\{u \in \mathbb{R}^m_{\geq 0} : u_1 = u_2 \geq \cdots \geq u_m = 0\right\}. \tag{5}$$

Theorem 4.6. *For any stakes function s and proportional recomposition, we have that*

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

Proof. 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} from (1), such that $\max_{a\in[m]} u^{\text{upper}}(a) = 1$. Similarly, let u^{lower} denote the utility vector in $\mathcal{U}$ that minimizes the infimum in (5). 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

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

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 $\tilde{\kappa}^{\text{lower}}$, the stakes-weighted plurality score obtained by a' is $p\kappa^{\text{upper}}$. Any other alternative $a \neq a'$ obtains a stakes-weighted plurality score of $(1-p)\tilde{\kappa}^{\text{lower}}/(m-1)$. Thus, a' winning the election amounts to the inequality

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

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

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

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

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

□

C OMITTED PROOFS FROM SECTION 5

C.1 Proof of Theorem 5.1

THEOREM 5.1. *For all randomized voting rules f and all 1-homogeneous stakes functions s,*

$$dist^s(f) \geq \frac{\sqrt{m}}{10 + 3\log m}.$$

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:

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

Next, we also divide the population into $\sqrt{m}$ parts $H_1, \ldots, H_{\sqrt{m}}$ 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 $\mathrm{sw}(\bar{a}, U) = 1/\sqrt{m}$.

(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

$$\frac{s\left(\mathbf{1}_{\log(m)-1}\mathbf{0}_{m-\log(m)+1}\right)}{s(\mathbf{1}_1\mathbf{0}_{m-1})} \geq 2^{\log m - 2} \geq \frac{m}{e^2}. \tag{6}$$

Now let us consider a histogram profile where the population is divided in $\sqrt{m}$ many equal sizes groups, which first-rank alternatives $1, \ldots \sqrt{m}$, 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

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

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 (6), 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

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

and the proof is complete. □

C.2 Formalisms about the Stable Lottery Rule

We now define the Stable lottery rule,following [7]. 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 C.1 (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. [7]. Building on this definition, we define the Stable Lottery Rule in terms of histograms.

Definition C.2 (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 Corollary 5.2. First, assume that $s \in \{\mathsf{max}, \mathsf{sum}\}$, and let f = Stable Lottery Rule. Then, by a well-established result from Ebadian *et al* [7], we know that both for $s = \mathsf{sum}$ and $s = \mathsf{max}$, 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 3.1 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 A.1. Indeed, this is seen as follows. Let $\mathbf{h}$ be an arbitrary preference histogram. In [7], 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 $s \in \{\mathsf{max}, \mathsf{sum}\}$.

It remains to show the claim for $s = \mathsf{range}$. Here, we argue along the same lines as Observation B.2: The worst-case distortion both for $s = \mathsf{range}$ and for $s = \mathsf{max}$ 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,

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

□

C.3 Folklore: all randomized rules have at least m distortion.

Fact C.3. *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 $\mathbb{E}[\mathrm{sw}(a', U)/n] \leq 1/m^2$, while the maximum welfare alternative has average utility $\mathrm{sw}(a, U)/n = 1/m$; thus the distortion of f is at least m. □

D OMITTED PROOFS FROM SECTION 6

D.1 Proof of Theorem 6.1

Theorem 6.1. *Let f be any voting rule and let s be any 1-homogeneous stakes function. Then, for any $\delta \geq 1$,*

$$dist^{\delta,s}(f) \leq \delta \cdot dist^{s}(f).$$

Proof. Fix a utility matrix U, a stakes function s and an error vector $\boldsymbol{\delta}$. Then, let $\tilde{U}$ be the utility matrix where voter i's utility vector is scaled by a factor 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

$$\mathsf{hist}^{s}(\tilde{U}) = \mathsf{hist}^{\boldsymbol{\delta},s}(U).$$

This directly implies that $f(\mathsf{hist}^{s}(\tilde{U})) = f(\mathsf{hist}^{\boldsymbol{\delta},s}(U))$. 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 in turn 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. □

D.2 Proof of Theorem 6.2

Theorem 6.2. *Let s be a 1-homogeneous stakes function such that $s(\mathbf{1}_2, \mathbf{0}_{m-2})$ is positive. Then, there exists a U such that $dist^{s}_{U}(\text{Plurality}) \geq (m-1) \cdot dist_{U}(\text{Plurality})$.*

Proof. Fix an arbitrary stakes function s that satisfies the conditions in the theorem. In this proof, we will denote an alternative a's plurality score in histogram $\mathbf{h}$ as $\text{pscore}(\mathbf{h}, a)$.

Step 1: Constructing core lower bound. We will construct a histogram $\mathbf{h}_{core}$ in which a' is the Plurality winner, and construct an underlying utility matrix U such that $\mathbf{h}_{core} = \mathsf{hist}(U)$. First, define a histogram $\mathbf{h}_{core}$ in which a' is ranked first by $1/(m-1)$ fraction of voters, and the other alternatives are each ranked first at equal frequency among the remaining voters. Let a^* (which will be the highest-welfare alternative in U) be ranked second when it is not ranked first, and let a' be ranked last when it is not ranked first. a' is indeed the winner, as it is ranked first more than any other alternative – moreover, for any other alternative $a \neq a'$,

$$\text{pscore}(\mathbf{h}_{core}, a') - \text{pscore}(\mathbf{h}_{core}, a) = \frac{1}{m-1} - \frac{m-2}{(m-1)^2} \geq \frac{1}{(m-1)^2}. \tag{7}$$

Now, realize $\mathbf{h}_{core}$ by giving all voters the ordered utility vector $\mathbf{u} = \mathbf{1}_2\mathbf{0}_{m-2}$. Then, $\mathsf{sw}(a^*, U)/n = 1$, and $\mathsf{sw}(a', U)/n = 1/(m-1)$. It follows that $\mathsf{dist}_U(\text{Plurality}) = m - 1$, which we will use later.

Step 2: Constructing instance-wise lower bound. Now, we want to construct a $\tilde{U}$ such that $\mathsf{dist}^{s}_{\tilde{U}}(\text{Plurality}) \geq (m-1) \cdot \mathsf{dist}_{\tilde{U}}(\text{Plurality})$. In doing so, we will take convex combinations of two histograms, where for $\alpha \in [0, 1]$ the π-th entry of $\alpha\mathbf{h}' + (1-\alpha)\mathbf{h}''$ is equal to $\alpha h'_\pi + (1-\alpha)h''_\pi$.

The first histogram in our convex combination will be $\mathbf{h}_{core}$; the second will be the histogram $\mathbf{h}_{pad}$, composed entirely of rankings in which a^* is ranked first, a' is ranked last, and all other alternatives are ranked arbitrarily. Now, for any $1/4 \le \alpha < 1/2$, define our final histogram $\mathbf{h}$ as

$$\mathbf{h} := \alpha \mathbf{h}_{core} + (1-\alpha)\mathbf{h}_{pad}.$$

Now, realize $\mathbf{h}$ with the utility matrix $\tilde{U}$ as follows: first, let N_{core} be the voters in the α fraction of $\mathbf{h}$ originally from $\mathbf{h}_{core}$, and let N_{pad} be those originally from $\mathbf{h}_{pad}$. For voters in N_{core}, realize their utilities as before, so they all have ordered utility vectors $\mathbf{u} = \mathbf{1}_2\mathbf{0}_{m-2}$. Then, for some arbitrarily small $\gamma > 0$, realize the rankings of voters in N_{pad} with the ordered utility vector $\gamma/3 \cdot \mathbf{1}_2\mathbf{0}_{m-2}$. Note that by construction we have that $\mathbf{h} = \mathsf{hist}(\tilde{U})$.

Analysis of $dist_{\tilde{U}}$(Plurality). Since a^* was the highest-welfare alternative in U and we only increased its welfare in $\tilde{U}$ at least as much as any other alternative, it remains the highest-welfare alternative in $\tilde{U}$. Moreover, because $\alpha < 1/2$, we have that Plurality($\mathbf{h}$) = Plurality($\mathsf{hist}(\tilde{U})$) $= a^*$, since Plurality is majority consistent. Thus, since $\mathsf{hist}(\tilde{U}) = \mathbf{h}$ by construction,

$$\mathrm{dist}_{\tilde{U}}(\textsc{Plurality}) = 1.$$

Analysis of $dist^s_{\tilde{U}}$(Plurality). By the 1-homogeneity of s, $s(\gamma/3 \cdot \mathbf{1}_2\mathbf{0}_{m-2}) = \gamma/3 \cdot s(\mathbf{1}_2\mathbf{0}_{m-2})$. Then,

$$\frac{\sum_{i\in N_{pad}} s(\mathbf{u}_i)}{\sum_{i\in[n]} s(\mathbf{u}_i)} \le \frac{\sum_{i\in N_{pad}} \gamma/3 \cdot s(\mathbf{1}_2\mathbf{0}_{m-2})}{\sum_{i\in N_{core}} s(\mathbf{1}_2\mathbf{0}_{m-2})} = \frac{(1-\alpha)\gamma/3}{\alpha} \le \gamma$$

where the last step uses $\alpha \ge 1/4$. All voters in N_{core} had the same stakes, so their relative frequency doesn't change from $\mathbf{h}_{core}$ to $\mathbf{h}^s$; moreover, all the voters in N_{pad} compose at most a γ fraction of the s-proportional electorate. From these facts, it follows that its histogram $\mathbf{h}^s$ must be very similar to $\mathbf{h}_{core}$; formally, for all π, we have that $h^s_\pi \in h_{core,\pi} \pm \gamma$. It follows that for all $a \in [m]$, $|\mathrm{pscore}(\mathbf{h}^s, a) - \mathrm{pscore}(\mathbf{h}_{core}, a)| \le \gamma$. Choosing $\gamma < 1/(m-1)^2$, by Equation (7), we have that $f(\mathbf{h}_{core}) = a' \implies f(\mathbf{h}^s) = a'$. Since in the construction of $\tilde{U}$ we only increased a^*'s social welfare relative to a', we get the central inequality in this chain, and we conclude the claim:

$$m - 1 = \mathrm{dist}_U(\textsc{Plurality}) = \frac{\mathsf{sw}(a^*, U)}{\mathsf{sw}(a', U)} \le \frac{\mathsf{sw}(a^*, \tilde{U})}{\mathsf{sw}(a', \tilde{U})} = \mathrm{dist}^s_{\tilde{U}}(\textsc{Plurality}).$$

□

D.3 Formalisms for the affected minority condition

First, we formally define the *affected minority* condition:

Definition D.1 (affected minority). U contains an affected minority if there exists some subset of voters $N_{am} \subseteq [n]$ satisfying the following conditions. Here, we let $\overline{N}_{am} = [n] \setminus N_{am}$ denote the set of voters who are not in this minority.

(1) $|N_{am}|/n \le 1/(m+1)$
(2) For all $i \in N_{am}$, we have that $s(\mathbf{u}_i) \ge \frac{m^2}{|N_{am}|/n} \cdot \max_{i' \in \overline{N}_{am}} s(\mathbf{u}_{i'})$
(3) Let $\tilde{A}_{am}$ be the set of alternatives a such that there exists any voter $i' \in \overline{N}_{am}$ who ranks a first. Then, for all $a \in \tilde{A}_{am}$,

$$\frac{1}{|N_{am}|} \sum_{i \in N_{am}} u_i(a) \le \frac{1}{|\overline{N}_{am}|} \sum_{i' \in \overline{N}_{am}} u_{i'}(a)$$

Now, we prove the sufficiency of this condition:

Theorem D.2. *If U contains an affected minority, then $dist^{max}_U$(Plurality) $\le$ $dist_U$(Plurality).*

Proof. The high level approach will be to prove that the existence of an affected minority implies that the winner of the standard election must have social welfare at least m times lower than at least one candidate preferred by members of the affected minority (and thus also m times less than the highest welfare alternative). Thus, the distortion of the standard election must be at least m and it follows from Corollary 4.5 that $\text{dist}^s_U(f) \leq \text{dist}_U(f)$. Let U be a utility matrix that contains an affected minority N_{am} (let $\overline{N}_{am} = [n] \setminus N_{am}$). Let a^* be the highest-welfare alternative in U, and let $a' = f(\text{hist}(U))$. Let A_{am} be the set of alternatives that are ranked first by any voter in N_{am}. For shorthand, let $v = \max_{i \in \overline{N}_{am}} s(\mathbf{u}_{i'})$ be the highest stakes of any member of $\overline{N}_{am}$.

First, we lower-bound $\text{sw}(a^*, U)$ by the observation that there must be an alternative $a \in A_{am}$ that is ranked first by at least $1/m$ fraction of voters in N_{am}. Then, by $s = \max$ and part (2) of Definition D.1, these voters must have utility $u_i(a) = s(\mathbf{u}_i) \geq \frac{m^2}{|N_{am}|/n} \cdot v$. Using that they compose at least a $|N_{am}|/(nm)$ fraction of all voters,

$$\text{sw}(a^*)/n \geq \text{sw}(a)/n \geq |N_{am}|/(nm) \cdot \frac{m^2}{|N_{am}|/n} \cdot v = mv$$

Now, we upper bound the social welfare of a'. Using part (1) of Definition D.1, we have that $|N_{am}|/n \leq 1/(m+1) < 1/m$, meaning that $a' \in \tilde{A}_{am}$, since it must be ranked first by at least one person in $\overline{N}_{am}$ to get a $1/m$ fraction of the votes. This voter can have at most utility v for this alternative, and moreover, by definition, the average utility for a' among voters in $\overline{N}_{am}$ can be at most v; then, by part (3) of Definition D.1, the average utility for a' among voters in N_{am} is also at most v. We conclude that $\text{sw}(a')/n \leq v$.

Combining these two bounds, we get that

$$\text{dist}_U(\textsc{Plurality}) = \frac{\text{sw}(a^*)}{\text{sw}(a')} \geq \frac{mv}{v} = m.$$

By Corollary 4.5, $\text{dist}^{\max}_U(\textsc{Plurality}) \leq m$, and the claim follows. □