

Improving Randomized Metric Distortion to 2.1441*

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

Abstract

In metric social choice, each voter ranks a set of m candidates by her distance to them in an unknown metric space. The cost of a candidate is its average distance to the voters. A randomized voting rule must use only the rankings to choose a lottery over candidates. Its distortion is the worst-case ratio between the expected cost under the lottery it returns and the cost of the best candidate. **CRWW24** prove an upper bound of 2.753, establishing a constant separation from deterministic rules, for which the best achievable distortion is 3. Recently, **Fra26** and **Ye26** independently improve the bound to 2.5.

We break this barrier by defining, proving the existence of, and using a new family of rules: *random-size stable lotteries*. Let D be a random variable over the domain of positive integers. A random-size stable lottery RSL_D guarantees that the probability of a random voter preferring any fixed candidate c to her favorite of D i.i.d. draws from RSL_D is at most $\mathbb{E}[1/(D+1)]$, where the probability averages over the voter, the value of D , and the D draws from RSL_D . When $D = 1$ deterministically, this is precisely the well-known maximal lottery. More generally, when $D = k$ deterministically, this reduces to its stable k -lottery generalization, which **CRTW25** prove the existence of for every fixed positive integer k . Their minimax argument for a fixed k easily generalizes to a random D .

Our main contribution is to show how stability with respect to a random D can be used to bound distortion directly. Specifically, we show that, for an explicit D supported on 1, 2, 9, and 10, any random-size stable lottery RSL_D has distortion at most 2.1441. The proof combines the biased-metric characterization of distortion due to **CRWW24** with a new potential argument, and exactly verifies a rational distortion bound via nonnegativity testing of polynomials in the Bernstein basis. We also prove a lower bound of 2.136 on the distortion of any randomized voting rule that, for any fixed D , always outputs a random-size stable lottery RSL_D . This is strictly higher than the best known universal lower bound of approximately 2.1126 on all randomized voting rules [CR22].

All the proofs have been obtained using GPT-5.6-Sol with significant guidance from the author and verified by the author, who significantly expanded on the exposition and simplified arguments with the aid of GPT-5.6-Sol and Claude Opus 5.

*A previous version of the manuscript provided a weaker distortion bound of 2.3282. That version used an infinite-dimensional linear program and its dual to bound distortion. The present version replaces that approach with a tighter potential-based analysis. A fresh grid search points to the draw distribution for which a distortion bound of 2.1441 can be verified exactly. This version also proves a lower bound of 2.136 for every rule in the family of rules we consider. Hence, for the particular rule used in our upper bound, the potential method's guarantee is less than 0.009 above the rule's exact distortion.

Contents

1	Introduction	3
1.1	Our Results	4
1.2	Technical Overview	5
1.3	Related Work	7
2	Preliminaries	8
3	Random-Size Stable Lotteries	8
4	Bounding the Metric Distortion via a Potential Function	11
4.1	Biased-metric characterization	11
4.2	Ranking coordinates and win probabilities	11
4.3	Bounding the offset loss	13
5	A Rule With Metric Distortion At Most 2.1441	18
6	A Lower Bound for Every Integer Draw Distribution	20
7	Discussion	21
A	Map of the Exact Upper-Bound Certificate	23
B	Further Properties of Random-Size Stable Lotteries	23
C	Proof of the Lower Bound	29

1 Introduction

Social choice theory studies how to choose one of several alternative decisions (candidates) by aggregating the conflicting preferences of multiple parties (voters). The norm is to elicit ranked preferences from voters over the candidates because exact utilities or costs are often difficult for voters to identify and report. Consequently, much of the centuries-old literature has sought voting rules that satisfy qualitative axioms defined on this ordinal data. However, no single axiom has been convincing on its own and seeking multiple axioms together has often led to impossibility results rather than a choice of the “optimal voting rule” [Arr51; Gib73; Sat75]. The rise of approximation algorithms in computer science has offered a new perspective: even when voters report ordinal preferences, one can aim to optimize welfare defined on the underlying numerical utilities or costs, with success measured by *distortion*—the approximation ratio between the achieved welfare and the optimum in the worst case over the unknown preference intensities. This has reignited the search for optimal voting rules, well-defined in this distortion framework.

Specifically, in the metric social choice model [ABEPS18], n voters and m candidates are points in the same underlying metric space. A voter’s distance to a candidate is her cost for that candidate, and she ranks candidates from smallest to largest distance. The social cost of a candidate is its average distance to the voters. A deterministic (resp., randomized) voting rule must choose a candidate (resp., a lottery over candidates) based only on the rankings, and its *distortion* is the ratio of the (expected) social cost of its choice to the smallest social cost of any candidate, in the worst case over all metrics consistent with the observed rankings.

GHS20 prove that the best distortion achievable by deterministic rules is 3. To see why deterministic rules cannot do any better and why randomization *may* help, consider a simple example with two voters (1, 2) and two candidates (a, b). Voter 1 prefers a to b while voter 2 prefers b to a . Let the underlying metric have distance d and $\Delta := d(a, b)$. Only two consistent possibilities are (provably) of interest:

- either voter 1 is located at a and voter 2 is located at the midpoint of a and b , which makes the social costs of a and b equal to $\Delta/4$ and $3\Delta/4$, respectively,
- or voter 1 is located at the midpoint and voter 2 is located at b , which makes the social costs of a and b equal to $3\Delta/4$ and $\Delta/4$, respectively.




Figure 1: Two metrics consistent with the reported rankings, which show a lower bound of 3 for deterministic rules and 2 for randomized rules.

Thus, as Figure 1 makes visible, either candidate can be optimal while the other is 3 times worse, making 3 a barrier that deterministic rules cannot cross. However, choosing between a and b uniformly at random achieves an expected social cost of $\Delta/2$, which is only 2 times the minimum

possible cost and that is the best possible distortion on this instance. It was even conjectured that 2 may be the optimal metric distortion of randomized rules, but **PS21**; **CR22** disprove it by establishing a lower bound strictly greater than 2, with the bound by the latter converging to approximately 2.1126 as $m \rightarrow \infty$.

On the upper-bound side, **CRWW24** prove a distortion guarantee of 2.753 by combining *maximal lottery*, the Nash equilibrium of a natural zero-sum game dating back to the 1960s [**Kre65**; **Fis84**], and RaDiUS, a novel rule they design. Their work, which received the SODA 2024 Best Paper Award, establishes that randomized rules can in fact achieve a constant distortion strictly lower than the best guarantee of 3 for deterministic rules. Recently, **Fra26** and **Ye26** independently improve this to 2.5 with a rule of striking simplicity: the equal mixture of a maximal lottery and a novel rule termed *Integrated Veto* (also, *Veto Lottery*), a randomized variant of the optimal deterministic rule *Plurality Veto* [**KK22**].¹ This still leaves a fundamental question open:

What is the smallest distortion that a randomized voting rule can achieve using only voters' preference rankings?

1.1 Our Results

The main result of our work improves the upper bound, closing about 92% of the gap between the upper bound of 2.5 and the lower bound of 2.1126 left by prior work.

Theorem 1. *There exists a randomized voting rule with metric distortion at most*

$$\frac{107200742653}{50000000000} = 2.14401485306.$$

To accomplish this, we consider a generalization of maximal lotteries to stable k -lotteries (SL_k) by **CRTW25**, who use them to design so-called tournament voting rules with low metric distortion; stable 1-lotteries are precisely maximal lotteries. We design a further generalization of this to *random-size stable lotteries*: given a *draw distribution* D over positive integers, its corresponding random-size stable lottery RSL_D effectively provides the guarantee of SL_k but with k drawn from D . Yet, we show that RSL_D cannot in general be obtained by simply mixing SL_k for different values of k ; we prove its existence for any D directly via a minimax argument similar to what **CRTW25** used for SL_k , and give a multiplicative-weights approximation algorithm with an explicit running time when D has finite support. The main technical contribution of our work is to connect the guarantee provided by RSL_D to metric distortion through a novel potential function and an elegant telescoping argument applied to the biased-metric characterization of distortion by **CRWW24**. The final distortion expression lends itself to a computer-assisted search for a suitable draw distribution D^* . This is followed by an exact verification using rational arithmetic, which reduces the problem to nonnegativity testing of polynomials using the so-called Bernstein basis.

While the upper bound of **Theorem 1** is eerily close to the best-known universal lower bound of approximately 2.1126 by **CR22**, our next result shows that the random-size stable lottery rule corresponding to any draw distribution D , which always returns an RSL_D lottery, has distortion more than 2.136, proving that this remaining gap cannot be closed simply by choosing a different draw distribution D . Note that the lower bound holds only when D is fixed independently of the profile.

¹Plurality Veto is a refinement of the Plurality Matching rule of **GHS20**, both achieving the optimal metric distortion of 3 for deterministic rules.

Theorem 2. *For any positive-integer-valued random variable D , there exists a profile on which every random-size stable lottery RSL_D has metric distortion strictly greater than*

$$\frac{267}{125} = 2.136.$$

Hence, this lower bound applies to the metric distortion of any randomized voting rule that always returns a random-size stable lottery RSL_D corresponding to any fixed D .

Together, [Theorems 1](#) and [2](#) place the best distortion achievable by the random-size stable lottery rule corresponding to any draw distribution D in the interval $[2.136, 2.14401485306]$. [Table 1](#) shows our upper and lower bounds in the context of prior work.

<i>Upper bounds</i>		
3	optimal deterministic rule	[GHS20]
2.753	a mixture of maximal lottery and RaDiUS	[CRWW24]
2.5	the equal mixture of maximal lottery and Integrated Veto	[Fra26; Ye26]
2.1441	a random-size stable lottery rule	this paper
<i>Lower bounds</i>		
2.0631	an example with $m = 7$	[PS21]
2.1126	a construction for any $m \geq 3$, bound on the left as $m \rightarrow \infty$	[CR22]
2.136	any random-size stable lottery rule	this paper

Table 1: Randomized metric distortion bounds from our work and prior work.

1.2 Technical Overview

Random-size stable lotteries. Our new ingredient starts from stable k -lotteries (SL_k) [CRTW25], which generalize maximal lotteries and satisfy

$$\Pr_v[c \succ_v (\text{SL}_k)^k] \leq \frac{1}{k+1} \quad \text{for every } c \in C.$$

Here $(\text{SL}_k)^k$ consists of k independent draws from SL_k , v is a uniformly random voter, and the probability also includes the draws and the prescribed tie-breaking. We let the number of draws itself be random. For a positive-integer-valued random variable D , a random-size stable lottery RSL_D satisfies

$$\Pr_v[c \succ_v (\text{RSL}_D)^D] \leq \mathbb{E} \left[\frac{1}{D+1} \right] \quad \text{for every } c \in C.$$

When $D = k$ deterministically, this recovers SL_k . Crucially, one lottery must be chosen before D is realized: independently choosing an SL_k for each k and then mixing those lotteries does not work because repeated draws from the mixture create uncontrolled cross terms. A convex minimax argument nevertheless proves that RSL_D exists for every D . The definition and proof appear in [Section 3](#); an approximation algorithm is deferred to [Appendix B](#).

Biased-metric characterization. The exact biased-metric characterization of **CRWW24**, adapted from an earlier linear programming approach of **CR22**, replaces the worst case over all metrics consistent with the observed ranked preference profile by a vector of m nonnegative offsets, with the offset of a designated optimal candidate normalized to zero. For each voter, the largest decrease between the offsets of a more-preferred and a less-preferred candidate determines her contribution to the optimal candidate’s cost, while the smallest offset among a candidate and all candidates ranked below it determines her contribution to that candidate’s excess cost over the optimum. Bounding distortion therefore reduces to comparing the lottery’s aggregate excess cost with the aggregate largest decrease for every offset vector.

Why a random number of draws helps. Fix a voter and a reference candidate c , and suppose that RSL_D places mass t on candidates strictly below c and no mass on c itself. The probability that c beats all D draws is

$$\phi(t) := \mathbb{E}[t^D] = \sum_{d \geq 1} \Pr[D = d] t^d.$$

It is tempting to use the same formula without the zero-mass assumption, but this is incorrect because draws equal to the reference candidate create ties. Under the random-label tie-breaking defined in [Section 3](#), if b is the mass strictly below the reference candidate and z is the mass on it, then conditioning on its label $s \in [0, 1]$ gives the corrected probability $\int_0^1 \phi(b + sz) ds$, as formalized in [Lemma 1](#). Averaging this over the voters gives a value at most $\mathbb{E}[1/(D + 1)]$, with equality when RSL_D puts a positive probability mass on c . The draw distribution determines the shape of ϕ , and therefore how strongly stability constrains different positions in the voters’ rankings.

Coming back to the offset formulation via biased metrics, for each value t , consider the candidates selected with positive probability whose offsets exceed t . These are the only ones that can contribute to the lottery’s excess cost at that value of t . The tightness established above limits how poorly these candidates can be positioned across the voters’ rankings on average, and thereby bounds their contribution to the excess cost. We express this bound as the difference of a one-dimensional potential W at two ranking percentiles. As t increases, the sets of candidates with offsets above t are nested, and the endpoint used at one level becomes the starting point at the next. The potential differences therefore telescope, leaving only $W(1)$. The biased-metric characterization then gives distortion at most $1 + 2W(1)$. Randomizing D lets us shape ϕ so that a single potential provides the required bound for every t . Computer-assisted search finds a D^* supported on $\{1, 2, 9, 10\}$, for which we construct a potential W satisfying the required inequalities with $W(1) = 57200742653/100000000000$. The resulting distortion bound of $107200742653/50000000000$ can be verified via exact rational arithmetic.

A lower bound for every integer draw distribution. We complement the potential upper bound with two explicit constructions of hard preference profiles. Each construction turns one scalar inequality involving the probability-generating function of D into an actual distortion witness. We use one instance of the first construction and two instances of the second. Except when $D = 1$ deterministically (handled separately), strict convexity makes the random-size stable lottery unique, so the construction applies to every way of selecting a stable lottery; a compactness argument converts weighted voters into a finite unweighted election. Finally, we show that no choice of the draw distribution can make all three resulting instances have distortion at most $267/125$.

1.3 Related Work

Metric distortion. The metric model was introduced by **ABEPS18**, who proved a lower bound of 3 for every deterministic rule and an upper bound of 5 for the Copeland rule. **GKM17** showed that Ranked Pairs can have distortion at least 5 and, more broadly, that even randomization cannot take tournament rules below distortion 3. **MW19** reduced the upper bound to $2 + \sqrt{5} \approx 4.236$ using generalized uncovered sets and proposed a candidate-selection condition whose universal existence would imply distortion 3; **Kem20LP** subsequently developed an LP-duality framework that unified these analyses and led to closely related sufficient conditions. **GHS20** finally matched the lower bound with Plurality Matching, settling the optimal deterministic distortion at 3; **KK22**; **KK23** subsequently obtained the same optimal guarantee with Plurality Veto and Simultaneous Plurality Veto. For randomized rules, **CRWW24** combine maximal lottery and a rule they call RaDiUS to obtain distortion 2.753. **Fra26** and **Ye26** independently improve the bound to 2.5 using the equal mixture of maximal lottery and Integrated Veto, called Veto Lottery by **Ye26**. Our random-size stable lottery improves this to 2.1441.

Utilitarian distortion. In the utilitarian model initiated by **PR06**, voters again report only rankings, but the hidden cardinal values are otherwise unrestricted and the objective is to maximize their sum. Here normalization is essential: rankings are unchanged if one rescales each voter’s utilities independently, so rank data alone provide no way to place different voters’ utilities on a common scale. The standard unit-sum model normalizes each voter’s utilities to sum to one, giving every voter equal weight while retaining the relative intensity of her preferences. Metric distortion needs no analogous normalization because a common underlying metric already places all voters’ costs on the same scale. Under unit-sum utilities, the optimal deterministic distortion is $\Theta(m^2)$ [**CP11**; **CNP+17**]; **BCHL+15** proved an $\Omega(\sqrt{m})$ lower bound for randomized rules, and **EKPS24** matched it with an $O(\sqrt{m})$ rule, closing the randomized case asymptotically as well. Although the metric and utilitarian models are each motivated by the same loss of cardinal information, rules optimized for one can perform poorly in the other. **GLS23** nevertheless design deterministic and randomized rules that simultaneously achieve near-optimal guarantees in both models, establishing best-of-both-worlds distortion.

Information–distortion tradeoffs. Full rankings are only one point on a broader elicitation spectrum. At one end, top- t ballots ask each voter to rank only her t favorite candidates; work under both utilitarian and metric objectives quantifies what can still be guaranteed as this reported prefix shrinks [**Kem20Comm**; **BHLS22**]. At the other end, a value query asks a voter for her exact utility for one candidate in addition to her ranking. A few such queries can sharply reduce utilitarian distortion [**ABFV21**]; for every fixed number of value queries per voter, **ES25** achieve distortion polynomial in $\min\{n, m\}$, with exponent equal to the reciprocal of the number of queries, matching the known dependence on m when $n = \Omega(m)$. More generally, one can optimize the elicitation format itself subject only to a bit budget. **MPSW19**; **MSW20** characterize this communication–distortion frontier up to logarithmic factors: achieving utilitarian distortion d requires $\tilde{\Theta}(m/d)$ bits per voter with deterministic elicitation and $\tilde{\Theta}(m/d^3)$ bits with randomized elicitation. Hence, the format of the elicited information can matter as much as its quantity: a carefully designed short message may yield lower distortion than a substantially longer ranked ballot.

2 Preliminaries

Let C be a finite set of m candidates and V a finite set of n voters. For a positive integer r , write $[r] = \{1, \dots, r\}$. Each voter $v \in V$ reports a strict ranking $\succ_v$ of C ; write $i \succsim_v j$ if $i = j$ or $i \succ_v j$. The *profile* is the collection $\sigma = (\succ_v)_{v \in V}$. Unless stated otherwise, v is drawn uniformly at random from V in every expectation or probability carrying the subscript v .

A *metric instance* places the voters and candidates in a common metric space with distance d .² It is *consistent* with σ if every voter's ranking agrees with her distances, that is, $a \succ_v b$ implies $d(v, a) \leq d(v, b)$ for all $v \in V$ and $a, b \in C$. The social cost of a candidate c and the expected social cost of a lottery $P \in \Delta(C)$, the set of probability distributions on C , are

$$SC_d(c) = \mathbb{E}_v[d(v, c)] \quad \text{and} \quad SC_d(P) = \mathbb{E}_{c \sim P}[SC_d(c)].$$

Definition 1: Distortion

For a profile σ and a lottery P , define

$$\text{dist}_\sigma(P) = \sup_{d \text{ consistent with } \sigma} \frac{SC_d(P)}{\min_{c \in C} SC_d(c)},$$

with $0/0 = 1$ and $z/0 = +\infty$ for $z > 0$. We drop the subscript when the profile is fixed.

A randomized rule $\mathcal{R}$ maps each profile σ to a lottery, and its *distortion* is $\sup_\sigma \text{dist}_\sigma(\mathcal{R}(\sigma))$. Note that the rule is unaware of the underlying metric.

3 Random-Size Stable Lotteries

We now formally define our key novel ingredient: random-size stable lotteries. We prove their existence by extending the minimax argument of **CRTW25** from a fixed number of draws k to a random number of draws D . In [Appendix B](#), we show that this also yields a standard multiplicative-weights algorithm to compute such lotteries approximately and establish other properties of interest that we will not use in our work.

Comparing multisets. First, we need to lift each voter's comparison of individual candidates to a comparison between multisets of candidates. Fix a voter v and two nonempty multisets S and T . It is natural to compare S and T by the candidates that v ranks highest in them. If these candidates are distinct, v 's ranking determines the comparison; the only ambiguity arises when the same candidate is highest-ranked in both multisets, possibly with different multiplicities. To resolve this ambiguity, we use the random-label tie-breaking scheme of **CRTW25**: give every duplicate an independent label drawn uniformly from $[0, 1]$, and let v compare duplicates by their labels. Thus, if voter v 's highest-ranked candidate in $S \cup T$ has multiplicity r in S and s in T , then $\Pr[S \succ_v T] = r/(r + s)$. By choosing S or T to be a singleton, this also yields comparisons between individual candidates and multisets.

Random sample size. Let D be a positive-integer-valued random variable indicating the number of draws. For a lottery P and a positive integer d , P^d and P^D denote the multisets obtained by d

²Distinct points are allowed to be at distance zero, so technically this is a pseudometric.

and D i.i.d. draws from P , respectively. To analyze the comparison between a candidate c and the multiset P^D , it will be helpful to introduce the following well-known quantity, which is an equivalent representation of a random variable.

Definition 2: Probability Generating Function

The probability generating function ϕ of D is

$$\phi(t) = \mathbb{E}[t^D] = \sum_{d \geq 1} \Pr[D = d] \cdot t^d.$$

For any D , ϕ is increasing and convex on $[0, 1]$, with $\phi(0) = 0$, $\phi(1) = 1$, and $\int_0^1 \phi(t) dt = \mathbb{E}[1/(D+1)]$. Unless $D = 1$ deterministically, ϕ is strictly convex.

The following lemma uses this function to compare a candidate c to the multiset P^D .

Lemma 1. Fix any voter v , lottery P , and candidate c . Then,

$$\Pr[c \succ_v P^D] = \int_0^1 \phi(b + sz) ds,$$

where $b = P(\{j : c \succ_v j\})$ is the P -mass on candidates v ranks strictly below c and $z = P(c)$. When $z > 0$, the same probability is $\frac{1}{z} \int_b^{b+z} \phi(u) du$. This probability is continuous and convex in P .

Proof. Let us first analyze the probability that c beats a random draw from P . Let $s \in [0, 1]$ be the random label assigned to c . Then, c beats all candidates ranked strictly below it (probability mass b) and an additional s fraction of its own probability mass which would be assigned a label smaller than s (probability mass $s \cdot z$). Hence, the total probability of c defeating a single draw from P is $b + sz$. The probability of it defeating d i.i.d. draws from P is then $(b + sz)^d$. Averaging over the number of draws d sampled from D and the random label s uniform in $[0, 1]$, we get

$$\begin{aligned} \Pr[c \succ_v P^D] &= \mathbb{E}_{d \sim D} \Pr[c \succ_v P^d] = \mathbb{E}_{d \sim D} \int_0^1 (b + sz)^d ds \\ &= \int_0^1 \left(\mathbb{E}_{d \sim D} (b + sz)^d \right) ds = \int_0^1 \phi(b + sz) ds, \end{aligned}$$

where the penultimate step uses Fubini's theorem to exchange the integral and expectation. When $z > 0$, the alternative expression follows by the variable substitution $u := b + sz$.

For each fixed d and s , the map $P \mapsto (b + sz)^d$ is continuous and convex because $b + sz$ is linear in P and $u \mapsto u^d$ is continuous and convex. Averaging over d and s preserves convexity. For continuity, if $P_k \rightarrow P$, the corresponding integrands converge pointwise and lie in $[0, 1]$, so dominated convergence over d and s applies. $\square$

Definition 3: Random-size stable lottery

We say that a lottery is *random-size stable* for D , and denote such a lottery by RSL_D , if

$$\forall c \in C : \Pr_v[c \succ_v (\text{RSL}_D)^D] \leq \mathbb{E} \left[\frac{1}{D+1} \right]. \quad (3.1)$$

If we have a challenger lottery A instead of a fixed challenger c , we can also equivalently write, denoting $A = A^1$ to be a random draw from A ,

$$\Pr_v[A \succ_v (\text{RSL}_D)^D] \leq \mathbb{E} \left[\frac{1}{D+1} \right].$$

This is at least as strong as the condition in [Definition 3](#), and equivalence follows from noticing that the left hand side is linear in A and, hence, maximized at some fixed c .

When $D = 1$ deterministically, [Definition 3](#) is precisely a maximal lottery; when $D = k$ deterministically, it is a stable k -lottery defined by [CRTW25](#), who prove its existence using a minimax argument.³ The next result shows that the same argument extends to a random number of draws D , proving the existence of RSL_D for all D .

Theorem 3 (Existence of random-size stable lotteries). *For every finite preference profile and every positive-integer-valued random variable D , a random-size stable lottery RSL_D exists. Further, recall [Equation \(3.1\)](#) defining RSL_D :*

$$\forall c \in C : \Pr_v[c \succ_v (\text{RSL}_D)^D] \leq \mathbb{E} \left[\frac{1}{D+1} \right].$$

No smaller right-hand side works for every profile, and every candidate in the support of RSL_D satisfies this with equality.

Proof. Let us abbreviate $\mu = \mathbb{E}[1/(D+1)]$. For a challenger lottery A and a lottery P , define

$$F(A, P) = \Pr_v[A \succ_v P^D].$$

The function is linear in A . [Lemma 1](#) establishes its continuity and convexity in P . Von Neumann's minimax theorem therefore gives

$$\min_P \max_A F(A, P) = \max_A \min_P F(A, P).$$

For every lottery P , exchangeability among $d+1$ independent labeled draws gives $\Pr_v[P \succ_v P^d] = 1/(d+1)$, and hence $F(P, P) = \mu$. Choosing $P = A$ shows that the right-hand side of the minimax equality is at most μ , while $\max_A F(A, P) \geq F(P, P) = \mu$ for every P . The minimax value is therefore exactly μ . Because F is linear in A , an optimal P , denoted RSL_D , satisfies [Equation \(3.1\)](#) for every candidate. Finally,

$$\mu = F(\text{RSL}_D, \text{RSL}_D) = \sum_{c \in C} \text{RSL}_D(c) \Pr_v[c \succ_v (\text{RSL}_D)^D].$$

Every term in this weighted average is at most μ , so equality holds for every candidate in the support. The same identity shows that no smaller universal right-hand side is possible. $\square$

Random-size stable lotteries cannot in general be obtained by simply averaging stable k -lotteries for different values of k ; [Appendix B.1](#) gives a two-candidate example.

The use of the minimax theorem also yields a standard multiplicative-weights algorithm for approximating RSL_D ; the precise guarantee appears as [Proposition 1](#) in [Appendix B](#).

An interesting identity, in which we reverse which side receives the D draws, is given as [Lemma 8](#) in [Appendix B](#).

³[CJMW20](#); [JMW20](#) define a related notion of stable lotteries over committees, which are more restrictive in that the committee does not necessarily come from independent draws from the same lottery, but more general in other dimensions, as discussed in [Appendix B](#).

4 Bounding the Metric Distortion via a Potential Function

We now connect random-size stability directly to metric distortion. The argument uses the exact biased-metric characterization of **CRWW24**, adapted from **CR22**, which reduces the worst case over all consistent metrics to worst case over m -dimensional vectors of nonnegative offsets.

4.1 Biased-metric characterization

Definition 4: Offset Loss and Offset Scale

Given an offset vector $x \in \mathbb{R}_{\geq 0}^C$, define

$$s_{v,c}(x) = \min_{j: c \succ_v j} x_j, \quad \Delta_v(x) = \max_{i,j: i \succ_v j} (x_i - x_j).$$

For a lottery P and a voter v , let the *offset loss of voter v* and the (overall) *offset loss*, respectively, be

$$\mathcal{L}_v(P, x) = \sum_{c \in C} P(c) s_{v,c}(x), \quad \mathcal{L}(P, x) = \mathbb{E}_v[\mathcal{L}_v(P, x)].$$

The *offset scale* is $\mathbb{E}_v[\Delta_v(x)]$.

In the metric supplied by the characterization below, the offset loss is the excess social cost of P over the designated optimal candidate, while the offset scale is twice that candidate's social cost.

Theorem 4 (Reduction to offsets [**CR22**; **CRWW24**]). *Fix a profile and a lottery P . Then*

$$\text{dist}(P) = \sup_{x \in \mathbb{R}_{\geq 0}^C: \min_c x_c = 0} \left(1 + 2 \frac{\mathcal{L}(P, x)}{\mathbb{E}_v[\Delta_v(x)]} \right),$$

with $0/0 = 0$ and $z/0 = +\infty$ for $z > 0$ in the displayed fraction. Further, for each candidate o and offset vector $x \in \mathbb{R}_{\geq 0}^C$ with $x_o = 0$, there exists a metric d consistent with the profile such that

$$\text{SC}_d(o) = \frac{1}{2} \mathbb{E}_v[\Delta_v(x)] \quad \text{and} \quad \text{SC}_d(c) - \text{SC}_d(o) = \mathbb{E}_v[s_{v,c}(x)], \forall c,$$

under which o is socially optimal and $\text{SC}_d(P) / \text{SC}_d(o) = 1 + 2\mathcal{L}(P, x) / \mathbb{E}_v[\Delta_v(x)]$.

For a finite profile, or more generally any distribution over rankings, we call the expression inside the supremum in **Theorem 4** the *offset value* of (P, x) . On a finite profile, for every $\lambda \geq 0$, the lottery P has distortion at most $1 + 2\lambda$ if and only if $\mathcal{L}(P, x) \leq \lambda \mathbb{E}_v[\Delta_v(x)]$ for every nonnegative offset vector x with $\min_{c \in C} x_c = 0$.

4.2 Ranking coordinates and win probabilities

We next develop a consequence of random-size stability that will let us bound $\mathcal{L}(P, x)$. Recall the probability generating function ϕ of D from **Definition 2**. We now define its integral and total mass.




Figure 2: The P -mass arranged from voter v 's least- to most-preferred candidate.

Definition 5: Integral and Stability Benchmark

For a positive-integer-valued random variable D with probability generating function ϕ , define

$$\Phi(u) = \int_0^u \phi(s) ds, \quad \mu = \Phi(1) = \mathbb{E} \left[\frac{1}{D+1} \right].$$

Next, we propose a visual interpretation that will allow us to understand and use [Lemma 1](#), which expresses $\Pr[c \succ_v P^D]$ in terms of ϕ , at a much more granular level.

Definition 6: Rank intervals, win probabilities, and deviations

For a lottery P and a voter v , arrange the mass placed by P on the different candidates on a line, from the voter's least-preferred candidate at coordinate 0 to her most-preferred candidate at coordinate 1, as illustrated in [Figure 2](#).

Each candidate c occupies a block $I_{v,c}$ of length $P(c)$. For $H \subseteq \text{supp}(P)$, define

$$H_v = \bigcup_{c \in H} I_{v,c}, \quad M_v(H) = \int_{H_v} \phi(u) du, \quad \tilde{M}_v(H) = M_v(H) - P(H)\mu.$$

Thus, H_v is the part of voter v 's interval occupied by H . A labeled copy of candidate c corresponds to a uniformly random point $u \in I_{v,c}$. By [Lemma 1](#), this copy beats D i.i.d. labeled draws from P with probability $\phi(u)$. Hence, $M_v(H)$ is the probability that a challenger drawn from P lies in H and beats D i.i.d. draws from P in voter v 's ranking. We call $M_v(H)$ the *win probability* of H for voter v and $\tilde{M}_v(H)$ its *win-probability deviation* from the benchmark $P(H)\mu$.

The next lemma justifies calling $P(H)\mu$ the benchmark. It is indeed the average of $M_v(H)$, making $\tilde{M}_v(H)$ a mean-zero deviation.

Lemma 2. *Let P be random-size stable for D . Then, for every $H \subseteq \text{supp}(P)$, we have $\mathbb{E}_v[\tilde{M}_v(H)] = 0$.*

Proof. Fix a voter v and candidate $c \in \text{supp}(P)$. [Lemma 1](#) shows that

$$\Pr[c \succ_v P^D] = \frac{1}{P(c)} \int_{I_{v,c}} \phi(u) du.$$

The final statement of [Theorem 3](#) says that this probability, averaged over voters, is exactly μ for every $c \in \text{supp}(P)$. Multiplying by $P(c)$ and summing over $c \in H$ gives $\mathbb{E}_v[M_v(H)] = P(H)\mu$, which is equivalent to the claim. $\square$

For any voter v , $P(H)\mu$ is the natural benchmark for $M_v(H)$: $P(H)$ is the probability that the challenger belongs to H , while μ is the probability that a distinguished draw is best among $D + 1$ i.i.d. draws. Thus, $\tilde{M}_v(H)$ is positive when challengers from H win more often than this benchmark from voter v 's perspective, and negative when they win less often. [Lemma 2](#) says that these deviations cancel when averaged over voters.

4.3 Bounding the offset loss

[Theorem 4](#) shows that the distortion of any lottery P can be upper bounded by $1 + 2\lambda$ if we show $\mathcal{L}(P, x) = \mathbb{E}_v[\mathcal{L}_v(P, x)] \leq \lambda \cdot \mathbb{E}_v[\Delta_v(x)]$ for all offset vectors $x \in \mathbb{R}_{\geq 0}^C$ with $x_o = 0$ for at least one candidate o . Fix any designated optimal candidate o and nonnegative offset vector x with $x_o = 0$. Further, fix any threshold $t \geq 0$ and mark each candidate c with $x_c > t$.

Proof strategy. If we had a voter-level bound $\mathcal{L}_v(P, x) \leq \lambda \cdot \Delta_v(x)$ for each voter v , then the above inequality would follow immediately by averaging over v . However, this may not hold for a reasonable value of λ . Instead, we will shift each voter's loss by adding $\tilde{M}_v(H)$ with the same multiplier for every voter, and then bound the shifted loss by $\lambda \cdot \Delta_v(x)$. Because the shifts are mean-zero ([Lemma 2](#)), they cancel out when averaging over v , which would still imply the desired inequality.

A final piece of the puzzle is to choose H . For a fixed threshold t , we will let H contain the marked candidates in the support of P whose offsets exceed t ; integrating over all possible t in the end will account for the entire offset.

Another objective we will achieve along the way is to compress all useful information from the profile into three quantities, a lower cutoff $a_v(t)$, a marked probability $p(t)$, and an upper cutoff $A_v(t)$, as functions of t . This will be helpful in computer-assisted search later on. We will begin by expressing $\mathcal{L}_v(P, x)$ in terms of the lower cutoff $a_v(t)$, and later express the correction term involving $M_v(H)$ in terms of all three quantities.

Definition 7: Lower cutoff

For each voter v , consider the P -mass arranged from her least-preferred candidate to her most-preferred in [Figure 2](#). Scan this interval from 0 to 1 and stop upon reaching the first unmarked candidate. The *lower cutoff* $a_v(t)$ is the total probability under P of the candidates passed before the scan stops.

In [Figure 2](#), this scan certifies that the interval from 0 through $a_v(t)$ is occupied by marked candidates. Its connection to the offset loss is exact.

Lemma 3 (Decomposing the offset loss into lower cutoffs). *For every voter v and threshold $t \geq 0$,*

$$a_v(t) = \Pr_{c \sim P}[s_{v,c}(x) > t].$$

Consequently,

$$\begin{aligned} \mathcal{L}_v(P, x) &= \int_0^\infty a_v(t) dt, \\ \mathcal{L}(P, x) &= \int_0^\infty \mathbb{E}_v[a_v(t)] dt. \end{aligned}$$

Proof. Recall that $s_{v,c}(x) = \min_{j: c \succsim_v j} x_j$. Thus, $s_{v,c}(x) > t$ if and only if c and every candidate ranked below c by v are marked, which holds exactly when the scan passes c . This proves the formal expression for $a_v(t)$. The tail-integral identity for the nonnegative random variable $s_{v,c}(x)$ then gives

$$\mathcal{L}_v(P, x) = \mathbb{E}_{c \sim P}[s_{v,c}(x)] = \int_0^\infty \Pr_{c \sim P}[s_{v,c}(x) > t] dt = \int_0^\infty a_v(t) dt.$$

Averaging over voters proves the second equality. $\square$

Definition 8: Marked probability and upper cutoff

Let $H(t) = \{c \in \text{supp}(P) : x_c > t\}$ be the marked support set, and let $p(t) = P(H(t))$ be the total probability that P places on marked candidates. For voter v , let

$$H_v(t) = \bigcup_{c \in H(t)} I_{v,c}$$

be the part of her ranking interval occupied by these candidates. An *upper cutoff* $A_v(t)$ is any coordinate satisfying $H_v(t) \subseteq [0, A_v(t)]$.

The scan in [Definition 7](#) passes only marked candidates, so $[0, a_v(t)] \subseteq H_v(t)$. The marked region $H_v(t)$ has length $p(t)$ and, by the definition of the upper cutoff, lies inside $[0, A_v(t)]$. Hence,

$$[0, a_v(t)] \subseteq H_v(t) \subseteq [0, A_v(t)] \quad \text{and} \quad 0 \leq a_v(t) \leq p(t) \leq A_v(t) \leq 1.$$

The smallest upper cutoff ends at the block of voter v 's most-preferred marked candidate. We also allow larger choices because this flexibility will be helpful when comparing successive thresholds. [Figure 3](#) illustrates the two cutoffs and the marked probability.




Figure 3: The example continues [Figure 2](#). At the threshold shown, $H(t) = \{c_5, c_4\}$, so the marked probability is $p(t) = P(c_5) + P(c_4)$. The scan for voter v passes c_5 and stops at the unmarked candidate c_2 , giving the lower cutoff $a_v(t) = P(c_5)$. The displayed upper cutoff $A_v(t)$ is deliberately nonminimal: it lies to the right of both marked blocks. Panel (b) keeps the forced interval $[0, a]$ and moves the remaining marked length $p - a$ as far right as the upper cutoff A permits.

The next lemma allows us to forget the exact locations of the marked blocks by bounding the win probabilities $M_v(H(t))$ in terms of only three quantities: the lower cutoff $a_v(t)$, their total length $p(t)$, and the upper cutoff $A_v(t)$.

Lemma 4 (Bounding the win probabilities in (a, p, A)). Let $0 \leq a \leq p \leq A \leq 1$, and let $S \subseteq [0, 1]$ be a measurable set of length p such that $[0, a] \subseteq S \subseteq [0, A]$. Then,

$$\int_S \phi(u) du \leq \Phi(a) + \Phi(A) - \Phi(a + A - p).$$

In particular, for every voter v and threshold t ,

$$M_v(H(t)) \leq \Phi(a_v(t)) + \Phi(A_v(t)) - \Phi(a_v(t) + A_v(t) - p(t)). \quad (4.1)$$

Proof. The part of S outside $[0, a]$ has length $p - a$. Because ϕ is nondecreasing, its integral is largest when this remaining length is placed at the right end of $[0, A]$, namely on $[a + A - p, A]$, as in Figure 3(b). This gives

$$\int_S \phi(u) du \leq \int_0^a \phi(u) du + \int_{a+A-p}^A \phi(u) du = \Phi(a) + \Phi(A) - \Phi(a + A - p).$$

Applying this bound to $S = H_v(t)$, using $[0, a_v(t)] \subseteq H_v(t) \subseteq [0, A_v(t)]$, yields Equation (4.1). $\square$

Adding a mean-zero correction. The lower cutoff $a_v(t)$ is the contribution that we need to bound, while Lemma 4 bounds the win probability $M_v(H(t))$. The win-probability deviation

$$\tilde{M}_v(H(t)) = M_v(H(t)) - p(t)\mu$$

links the two: because $H(t)$ and $p(t)$ do not depend on the voter, Lemma 2 gives $\mathbb{E}_v[\tilde{M}_v(H(t))] = 0$. We may therefore add any voter-independent multiple of this deviation to $a_v(t)$ without changing its average.

Suppose $0 < p(t) < 1$. Then $p(t)\mu - \Phi(p(t)) > 0$ because the average of the nondecreasing, nonconstant function ϕ on $[0, p(t)]$ is smaller than its average on $[0, 1]$. Consider the reference placement in which the scan passes every candidate in the marked support set. Here $a_v(t) = p(t)$, $H_v(t) = [0, p(t)]$, and $M_v(H(t)) = \Phi(p(t))$. We choose the multiplier so that the corrected lower cutoff is zero in this reference placement.

Definition 9: Lower-cutoff correction

When $0 < p(t) < 1$, define the *lower-cutoff correction* and the *corrected lower cutoff*, respectively, by

$$\begin{aligned} \text{corr}_v(t) &= \frac{p(t)}{p(t)\mu - \Phi(p(t))} \tilde{M}_v(H(t)), \\ Q_v(t) &= a_v(t) + \text{corr}_v(t). \end{aligned} \quad (4.2)$$

The multiplier is positive. Thus, substituting the win-probability bound Equation (4.1) into Equation (4.2) bounds $Q_v(t)$ using only the three scalars retained from the ranking. We name this scalar bound next.

Definition 10: Scalar bounds for corrected lower cutoffs

For $0 < p < 1$ and $0 \leq a \leq p \leq A \leq 1$, define

$$\mathcal{E}_\phi(a, p, A) = a + \frac{p(\Phi(a) + \Phi(A) - \Phi(a + A - p) - p\mu)}{p\mu - \Phi(p)}, \quad (4.3)$$

and set $\mathcal{E}_\phi(0, 0, A) = 0$. For $r \in [0, 1]$, define

$$\mathcal{G}_\phi(r) = r - \frac{\phi(r) - \mu}{1 - \mu}.$$

The function $\mathcal{E}_\phi$ handles $p(t) < 1$. When $p(t) = 1$, the denominator in Equation (4.2) vanishes, and we use $\mathcal{G}_\phi$ instead. Indeed, $x_o = 0 \leq t$ makes o unmarked, and $p(t) = 1$ therefore implies $P(o) = 0$. If r_v is the probability assigned to candidates that voter v ranks below o , then the lower-cutoff scan stops no later than o , giving $a_v(t) \leq r_v$, while Lemma 1 shows that stability against o bounds the average of $\phi(r_v)$. The following lemma records both cases.

Lemma 5 (Correcting one lower cutoff). *Let P be random-size stable for D . Fix a candidate o , a nonnegative offset vector x with $x_o = 0$, and a threshold $t \geq 0$. Use the lower cutoff from Definition 7 and the quantities from Definition 8. If $p(t) = 0$, set $Q_v(t) = 0$. If $0 < p(t) < 1$, let $Q_v(t)$ be the corrected lower cutoff from Definition 9. Then, for every upper cutoff $A_v(t)$,*

$$\mathbb{E}_v[Q_v(t)] = \mathbb{E}_v[a_v(t)] \quad \text{and} \quad Q_v(t) \leq \mathcal{E}_\phi(a_v(t), p(t), A_v(t)).$$

If $p(t) = 1$, let r_v be the probability assigned to candidates that voter v ranks below o and set $Q_v(t) = \mathcal{G}_\phi(r_v)$. Then $\mathbb{E}_v[Q_v(t)] \geq \mathbb{E}_v[a_v(t)]$.

Proof. Abbreviate $p = p(t)$. If $p = 0$, then $a_v(t) = 0$ for every voter, so both claims are immediate. If $0 < p < 1$, Lemma 2 gives $\mathbb{E}_v[\text{corr}_v(t)] = 0$, proving the equality of averages. The pointwise bound follows by substituting Equation (4.1) into Equation (4.2).

Suppose $p = 1$. As observed above, $P(o) = 0$ and $a_v(t) \leq r_v$. By Lemma 1, stability against o gives $\mathbb{E}_v[\phi(r_v)] \leq \mu$. Therefore,

$$\mathbb{E}_v[Q_v(t)] = \mathbb{E}_v[r_v] + \frac{\mu - \mathbb{E}_v[\phi(r_v)]}{1 - \mu} \geq \mathbb{E}_v[r_v] \geq \mathbb{E}_v[a_v(t)].$$

□

The one-threshold reduction. At a fixed threshold, the exact positions of the marked candidates affect $M_v(H(t))$, but Lemma 5 removes this dependence. For the remaining argument, we retain only the lower cutoff $a_v(t)$, the marked probability $p(t)$, and an upper cutoff $A_v(t)$; $\mathcal{E}_\phi(a_v(t), p(t), A_v(t))$ bounds the corrected lower cutoff.

From one threshold to all thresholds. To combine the one-threshold bounds for a fixed voter, we compare thresholds separated by $\Delta_v(x)$. At each threshold after the first, the definition of $\Delta_v(x)$ ensures that every candidate still marked is ranked below the candidate that stopped the scan at the preceding threshold. Consequently, the preceding lower cutoff is an upper cutoff at the current threshold.

The remaining analytic goal. We seek a potential W such that $\mathcal{E}_\phi(a, p, A) \leq W(A) - W(a)$ and $\mathcal{G}_\phi(r) \leq W(1) - W(r)$. When the preceding lower cutoff is used as the current upper cutoff, these potential differences telescope across thresholds and leave only $W(1)$.

The next lemma completes the reduction from the profile-dependent metric problem to a one-dimensional potential. Its conclusion has the factor $1 + 2W(1)$ because the offset characterization identifies the optimal candidate's cost with one half of the offset scale.

Lemma 6 (Potential criterion). *Let P be random-size stable for D . Suppose there is a nondecreasing function $W : [0, 1] \rightarrow \mathbb{R}$ with $W(0) = 0$ such that*

$$\mathcal{E}_\phi(a, p, A) \leq W(A) - W(a) \quad \text{for every } 0 \leq a \leq p \leq A \leq 1 \text{ with } p < 1, \quad (4.4)$$

$$\mathcal{G}_\phi(r) \leq W(1) - W(r) \quad \text{for every } r \in [0, 1]. \quad (4.5)$$

Then P has metric distortion at most $1 + 2W(1)$.

Proof. Fix a nonnegative offset vector $(x_c)_{c \in C}$ and a designated candidate o with $x_o = 0$. For each voter v , abbreviate $\Delta_v = \Delta_v(x)$ as in [Theorem 4](#). For every threshold $t \geq 0$, let $(Q_v(t))_v$ be the corrected lower cutoffs supplied by [Lemma 5](#).

We first prove the deterministic inequality

$$\int_0^\infty Q_v(t) dt \leq W(1)\Delta_v. \quad (4.6)$$

If $\Delta_v = 0$, the offsets cannot increase as we move upward through voter v 's ranking. At every threshold, the scan therefore passes all candidates in the marked support set before it stops, so $H_v(t) = [0, p(t)]$ and $a_v(t) = p(t)$. Taking $A_v(t) = p(t)$ bounds every corrected lower cutoff with $p(t) < 1$ by $\mathcal{E}_\phi(p(t), p(t), p(t)) = 0$. If $p(t) = 1$, every supported candidate lies below o , so $r_v = 1$ and $\mathcal{G}_\phi(1) = 0$. This proves [Equation \(4.6\)](#) when $\Delta_v = 0$.

Suppose $\Delta_v > 0$. Every threshold $t \geq 0$ can be written uniquely as $t = s + k\Delta_v$, where $s \in [0, \Delta_v)$ and k is a nonnegative integer. Fix the base threshold s and define

$$L_c = |\{k \geq 0 : x_c > s + k\Delta_v\}|.$$

Thus, L_c counts how many of the successive thresholds $s, s + \Delta_v, s + 2\Delta_v, \dots$ are smaller than x_c . The definition of Δ_v implies $L_i - L_j \leq 1$ whenever $i \succ_v j$: a more-preferred candidate cannot remain marked for two more thresholds than a less-preferred candidate. Put $H_k = H(s + k\Delta_v)$, $H_{k,v} = H_v(s + k\Delta_v)$, $p_k = p(s + k\Delta_v)$, and $a_k = a_v(s + k\Delta_v)$, and set $a_{-1} = 1$.

If $p_k < 1$, then $H_{k,v} \subseteq [0, a_{k-1}]$. For $k = 0$, this is immediate from $a_{-1} = 1$. For $k \geq 1$, a violation would place some candidate with $L_c \geq k + 1$ above the candidate with $L_c \leq k - 1$ that stopped the preceding scan, contradicting $L_i - L_j \leq 1$. Hence, a_{k-1} is an upper cutoff at threshold $s + k\Delta_v$. Applying [Lemma 5](#) and [Equation \(4.4\)](#) with $a = a_k$ and $A = a_{k-1}$ bounds $Q_v(s + k\Delta_v)$ by $W(a_{k-1}) - W(a_k)$.

If $p_k = 1$ and $k \geq 1$, every supported candidate has $L_c \geq k + 1 \geq 2$, whereas $L_o = 0$. The inequality $L_i - L_j \leq 1$ forces every supported candidate below o , so $r_v = 1$ and $Q_v(s + k\Delta_v) = \mathcal{G}_\phi(1) = 0$. Thus, only $k = 0$ can give a nonzero all-support contribution. In that case, any candidate in H_1 has $L_c \geq 2$ and must lie below o , so $H_{1,v} \subseteq [0, r_v]$. Hence, r_v is an upper cutoff at the next threshold. Any intervening all-support contributions vanish; when the marked probability

first drops below one, its marked region still lies inside $[0, r_v]$, so the subsequent bounds from Equation (4.4) telescope from $W(r_v)$ to zero. Therefore, Equation (4.5) gives

$$\sum_{k \geq 0} Q_v(s + k\Delta_v) \leq \mathcal{G}_\phi(r_v) + W(r_v) \leq W(1).$$

If no threshold has marked probability one, the bounds from Equation (4.4) telescope to $W(1)$ because $a_{-1} = 1$, while $a_k = 0$ for all sufficiently large k . Integrating over the base threshold yields

$$\int_0^\infty Q_v(t) dt = \int_0^{\Delta_v} \sum_{k \geq 0} Q_v(s + k\Delta_v) ds \leq W(1)\Delta_v,$$

which proves Equation (4.6).

It remains to average over voters and integrate over thresholds. Lemma 5 gives $\mathbb{E}_v[Q_v(t)] = \mathbb{E}_v[a_v(t)]$ when $p(t) < 1$ and $\mathbb{E}_v[Q_v(t)] \geq \mathbb{E}_v[a_v(t)]$ when $p(t) = 1$. All functions are bounded step functions with finite support in t , so Fubini's theorem and Equation (4.6) give

$$\mathcal{L}(P, x) = \int_0^\infty \mathbb{E}_v[a_v(t)] dt \leq W(1) \mathbb{E}_v[\Delta_v]. \quad (4.7)$$

Because Equation (4.7) holds for every o and every nonnegative x with $x_o = 0$, Theorem 4 gives $\text{dist}(P) \leq 1 + 2W(1)$. $\square$

5 A Rule With Metric Distortion At Most 2.1441

We apply Lemma 6 to the exact draw distribution in Figure 4. A numerical search suggested its sparse support and a value near 2.14401485303, but only the exact rational bound in the figure is claimed and verified over the complete continuum.

Figure 4: Candidate draw distribution and distortion bound

Let D^* be supported on $\{1, 2, 9, 10\}$ with

$$\begin{aligned} \Pr[D^* = 1] &= \frac{462688493532}{10^{12}}, & \Pr[D^* = 2] &= \frac{194607598780}{10^{12}}, \\ \Pr[D^* = 9] &= \frac{156952022902}{10^{12}}, & \Pr[D^* = 10] &= \frac{185751884786}{10^{12}}. \end{aligned}$$

Its probability-generating function, its integral, and its stability benchmark are

$$\begin{aligned} \phi(u) &= \frac{462688493532u + 194607598780u^2}{10^{12}} \\ &\quad + \frac{156952022902u^9 + 185751884786u^{10}}{10^{12}}, \\ \Phi(u) &= \int_0^u \phi(s) ds, \quad \mu = \Phi(1) = \frac{54251205298963}{16500000000000}. \end{aligned}$$

The potential parameters and candidate distortion bound are

$$t_0 = \frac{508613}{1000000}, \quad C_0 = \frac{57200742653}{100000000000}, \quad 1 + 2C_0 = \frac{107200742653}{50000000000} = 2.14401485306.$$

Theorem 1. *There exists a randomized voting rule with metric distortion at most*

$$\frac{107200742653}{50000000000} = 2.14401485306.$$

Proof. Let D^* , ϕ , Φ , μ , t_0 , and C_0 be as specified in Figure 4. Define

$$W(u) = \begin{cases} 0, & 0 \leq u \leq t_0, \\ C_0 - u - \frac{\mu - \phi(u)}{1 - \mu}, & t_0 < u \leq 1. \end{cases} \quad (5.1)$$

By Lemma 6, it suffices to verify four claims: the endpoints of W , the boundary condition Equation (4.5), monotonicity of W , and the internal condition Equation (4.4). All sign checks below use exact rational arithmetic. The polynomial inequalities are verified using the Bernstein subdivision method of MN13.⁴ The accompanying repository⁵ contains the static checker `certificate_distortion.py`. It verifies the recorded polynomial identities, every subdivision, and the required coefficient signs without performing a search. Appendix A maps each check below to its name in the program.

Verification 1: Endpoints. The displayed weights are positive and sum to one. By definition, $W(0) = 0$; because $\phi(1) = 1$, Equation (5.1) gives $W(1) = C_0$.

Verification 2: The boundary condition. When $u > t_0$, substituting Equation (5.1) makes Equation (4.5) an equality. When $u \leq t_0$, the boundary condition is equivalent to

$$C_0 - u - \frac{\mu - \phi(u)}{1 - \mu} \geq 0 \quad \text{for every } u \in [0, t_0]. \quad (5.2)$$

The checker verifies this polynomial inequality throughout the interval. In particular, at $u = t_0$, it shows that the right limit of W is at least $W(t_0) = 0$.

Verification 3: Monotonicity. The first branch of W is constant. Exact rational arithmetic verifies $\phi'(t_0) > 1 - \mu$. Because ϕ is convex, for $u > t_0$ we have $W'(u) = -1 + \phi'(u)/(1 - \mu) > 0$. Together with the junction check from Verification 2, this proves that W is nondecreasing.

Verification 4: The internal condition. The case $p = 0$ is immediate from the definition of $\mathcal{E}_\phi$. Suppose first that $0 < p < 1$ and $A \leq t_0$. Put $d = p - a$ and $K = p\mu - \Phi(p) > 0$. Multiplying Equation (4.3) by K and using that $\Phi(x)/x$ is nondecreasing gives

$$\begin{aligned} K\mathcal{E}_\phi(a, p, A) &= p(\Phi(A) - \Phi(A - d)) + p\Phi(a) - a\Phi(p) - pd\mu \\ &\leq pd(\phi(A) - \mu) \leq 0. \end{aligned}$$

⁴They express a polynomial on the current box in the so-called Bernstein basis, whose coefficients bound its values throughout the box. If these bounds do not establish the desired sign, they bisect the box, update the Bernstein coefficients on the two subboxes, and repeat.

⁵<https://github.com/nisarg89/metric-distortion>

The last inequality uses the exact check $\phi(t_0) < \mu$ and the monotonicity of ϕ . Hence $\mathcal{E}_\phi(a, p, A) \leq 0 = W(A) - W(a)$ throughout the low region.

On the remaining region, clear the positive denominator in Equation (4.3). The required slack is

$$\begin{aligned} S(a, p, A) = & (W(A) - W(a) - a)(p\mu - \Phi(p)) \\ & - p(\Phi(a) + \Phi(A) - \Phi(a + A - p) - p\mu). \end{aligned} \quad (5.3)$$

The breakpoint t_0 leaves three cases outside the low region: $a \leq p \leq t_0 \leq A$, $a \leq t_0 \leq p \leq A$, and $t_0 \leq a \leq p \leq A$. After mapping each case to the unit cube and cancelling factors that are nonnegative throughout the relevant domain, the checker verifies that the resulting rational polynomial is nonnegative. Thus $S \geq 0$ in all three cases. The face $A = t_0$ is covered by the analytic low-region argument; on the face $a = t_0$, using the right branch for $W(a)$ only strengthens the inequality because the actual value is $W(t_0) = 0$.

All four claims of Lemma 6 now hold with $W(1) = C_0$, so every RSL_{D^*} has the candidate distortion bound in Figure 4. On every profile, such a lottery exists by Theorem 3; choosing any one of them defines a randomized rule with the claimed distortion. $\square$

How the distribution was found. The numerical discovery phase is separate from the exact proof:

- The code `search_distortion.py` enumerated all $\sum_{i=1}^4 \binom{20}{i} = 6195$ subsets of $\{1, \dots, 20\}$ of size at most four and, for each support, numerically optimized its weights against a finite-grid longest-path relaxation of the potential criterion.
- We reoptimized the best 100 supports on a finer grid and applied differential evolution followed by sequential least squares programming (SLSQP) to the best 20. This produced the support $\{1, 2, 9, 10\}$ and the numerical parameters that were then rationalized. These finite-grid calculations neither certify a continuum bound nor establish global optimality.
- The generic adaptive verifier `verify_distortion.py` constructed the fixed exact certificate described above. The theorem relies only on the rational parameters and the non-searching checker `certificate_distortion.py`.

6 A Lower Bound for Every Integer Draw Distribution

The preceding upper bound uses a carefully chosen distribution of the number of draws. We now ask how much the guarantee could improve by choosing a different positive-integer-valued distribution while retaining random-size stability. The answer is limited: no fixed draw distribution can achieve distortion at most 2.136.

Theorem 2. *For any positive-integer-valued random variable D , there exists a profile on which every random-size stable lottery RSL_D has metric distortion strictly greater than*

$$\frac{267}{125} = 2.136.$$

Hence, this lower bound applies to the metric distortion of any randomized voting rule that always returns a random-size stable lottery RSL_D corresponding to any fixed D .

The proof has two conceptual ingredients. First, we construct two symmetric families of preference profiles. In the first, one candidate is inserted at one of two positions among exchangeable candidates; in the second, a distinguished set of candidates occupies one of two pairs of percentile intervals. In each case, random-size stability pins down the uniform lottery, while the offset characterization turns its position in the rankings into an exact distortion ratio. Whenever D is not deterministically one, the probability-generating function $x \mapsto \mathbb{E}[x^D]$ is strictly convex, which makes the stable lottery unique and removes any dependence on the rule’s selection convention. The idealized voter probabilities are then approximated by a finite electorate without losing the strict inequality.

Second, we reweight the probability of choosing each possible number d of draws by $1/(d+1)$. After this reweighting, whether each of the three ratios exceeds $71/125$ can be expressed by the sign of the expectation of an explicit function of d . We exhibit positive weights for which the weighted sum of these three functions is strictly negative for every positive integer d . Its expectation is therefore strictly negative for every draw distribution, so at least one of the three profiles has excess-cost ratio greater than $71/125$, and hence distortion greater than $1 + 2 \cdot 71/125 = 267/125$. The complete constructions, the finite-profile transfer, and the exact calculation appear in [Appendix C](#).

7 Discussion

There are now two distinct gaps. Within the family of rules that always return a random-size stable lottery RSL_D for some fixed integer draw distribution D , our lower and upper bounds are less than 0.009 apart. Closing this gap may require a more refined search over draw distributions and potentials and possibly tightening our potential-based method of bounding the loss. A near-optimal draw distribution may suggest a candidate for the optimal draw distribution, which may then be provable analytically.

The second possible gap is the one between our restricted family and the broader family of all randomized voting rules. While our upper bound is less than 0.04 above the universal (asymptotic) lower bound of **CR22**, our lower bound implies that no draw distribution can completely close this gap. That said, it is possible that the bound of **CR22** is not tight and an optimal randomized rule in fact lies within our family. However, if it is tight, closing the remaining gap will require seeking a rule outside of our family.

AI Disclosure

All the mathematical proofs were derived by OpenAI Codex (GPT-5.6-Sol at Max effort) based on research directions, literature connections, proof and search strategies, and inspirations supplied by the author. The author has verified all mathematical details, presented simplified arguments and detailed expositions, often with the aid of GPT-5.6-Sol and Claude Opus 5, and retains full responsibility for all the content and any errors.

Acknowledgments

This work was supported by an NSERC Discovery Grant and an NSERC-CSE Research Communities Grant. Researchers funded through the NSERC-CSE Research Communities Grants do not

represent the Communications Security Establishment Canada or the Government of Canada. Any research, opinions or positions they produce as part of this initiative do not represent the official views of the Government of Canada.

Appendix

A Map of the Exact Upper-Bound Certificate

This section records how the mathematical checks in the proof of [Theorem 1](#) appear in `certificate_distortion.py`. Substituting the polynomials ϕ , Φ , and W from [Figure 4](#) and [Equation \(5.1\)](#) into the cited expressions, including the slack S defined in [Equation \(5.3\)](#), and expanding gives the program’s polynomials; those routine expansions are omitted.

Check in the proof	Checker routine	Polynomial in the checker
The probabilities are positive and sum to one, and the distortion bound is $1 + 2C_0$	<code>exact_parameter_checks</code>	<code>probability_mass_identity</code> and its reverse, <code>probability_0_positive</code> through <code>probability_3_positive</code> , and <code>distortion_identity</code> and its reverse
$\phi(t_0) < \mu$ and $\phi'(t_0) > 1 - \mu$	<code>threshold_checks</code>	<code>activation_below_mean</code> and <code>threshold_derivative</code>
Equation (5.2) on $[0, t_0]$	<code>boundary_inequality</code>	<code>boundary_polynomial</code>
Monotonicity of the second branch of W on $[t_0, 1]$	<code>potential_monotonicity</code>	<code>potential_derivative</code> , a direct polynomial check of the conclusion also obtained from convexity in Verification 3
$a \leq p \leq t_0 \leq A$	<code>internal_cell_001</code>	<code>cell_001</code>
$a \leq t_0 \leq p \leq A$	<code>internal_cell_011</code>	<code>cell_011</code>
$t_0 \leq a \leq p \leq A$	<code>internal_cell_111</code>	<code>cell_111</code>

For completeness, the three cell polynomials use variables $(x, y, z) \in [0, 1]^3$. The polynomial `cell_001` is the polynomial continuation of $S(a, p, A)/p$ after substituting $a = t_0xy$, $p = t_0y$, and $A = t_0 + (1 - t_0)z$. The polynomial `cell_011` is the polynomial continuation of $S(a, p, A)/(1 - p)$ after substituting $a = t_0x$, $p = t_0 + (1 - t_0)y$, and $A = p + (1 - p)z$. Finally, `cell_111` is the polynomial continuation of $S(a, p, A)/((1 - a)(1 - p))$ after substituting $a = t_0 + (1 - t_0)x$, $p = a + (1 - a)y$, and $A = p + (1 - p)z$. The cancelled factors are positive in the interiors of the corresponding cells, and the polynomial continuations cover their boundary faces.

B Further Properties of Random-Size Stable Lotteries

This appendix records further consequences of random-size stability. It relates the definition to deterministic challenger committees, comparisons using other numbers of samples, and lotteries with small support. These consequences do not subsume stable lotteries over committees [[CJMW20](#); [JMW20](#)]: here every committee consists of independent draws from one lottery and a voter compares committees by their favorite members, whereas that literature permits correlated committee lotteries and more general committee preferences.

B.1 Why separately chosen stable lotteries cannot simply be mixed

Suppose masses $3/5$ and $2/5$ of the voters rank $a \succ b$ and $b \succ a$, respectively. The lottery $P_1 = \delta_a$ is stable for one draw, while $P_2 = (4/5)\delta_a + (1/5)\delta_b$ is stable for two draws because

$$\begin{aligned}\Pr_v[a \succ_v P_2^2] &= \frac{3}{5} \int_0^1 \left(\frac{1}{5} + \frac{4}{5}s \right)^2 ds + \frac{2}{5} \int_0^1 \left(\frac{4}{5}s \right)^2 ds = \frac{1}{3}, \\ \Pr_v[b \succ_v P_2^2] &= \frac{3}{5} \int_0^1 \left(\frac{1}{5}s \right)^2 ds + \frac{2}{5} \int_0^1 \left(\frac{4}{5} + \frac{1}{5}s \right)^2 ds = \frac{1}{3}.\end{aligned}$$

If D is uniform on $\{1, 2\}$, their average $P = (P_1 + P_2)/2$ satisfies

$$\Pr_v[a \succ_v P] = \frac{51}{100}, \quad \Pr_v[a \succ_v P^2] = \frac{33}{100},$$

and hence $\Pr_v[a \succ_v P^D] = 21/50 > 5/12 = \mathbb{E}[1/(D+1)]$. The cross terms created by drawing repeatedly from P are not controlled by the separate guarantees for P_1 and P_2 . Selecting P_D after observing D would satisfy a different averaged statement, not the requirement that one base lottery be chosen before D .

B.2 Computing a random-size stable lottery

Proposition 1. Suppose D has finite support of size s , whose largest value is $d_{\max}$, and let $m = |C| \geq 2$. For every positive integer N , one can compute a lottery $\bar{P}$ satisfying

$$\max_{c \in C} \Pr_v[c \succ_v \bar{P}^D] \leq \mathbb{E} \left[\frac{1}{D+1} \right] + \sqrt{\frac{\log m}{2N}}.$$

In the real-arithmetic model, the algorithm uses $O(Nnms \log(d_{\max} + 1))$ arithmetic operations and exponential evaluations. Thus additive error at most ε is obtained in $O(nms \log(d_{\max} + 1) \log(m)/\varepsilon^2)$ time.

Proof. Start with the uniform lottery P_1 and set $\eta = \sqrt{8 \log(m)/N}$. At iteration t , put $g_{t,c} = \Pr_v[c \succ_v P_t^D]$ and update

$$P_{t+1}(c) = \frac{P_t(c) \exp(\eta g_{t,c})}{\sum_{b \in C} P_t(b) \exp(\eta g_{t,b})}.$$

The standard exponential-weights calculation for numbers in $[0, 1]$ gives

$$\max_{c \in C} \frac{1}{N} \sum_{t=1}^N g_{t,c} \leq \frac{1}{N} \sum_{t=1}^N \sum_{c \in C} P_t(c) g_{t,c} + \frac{\log m}{\eta N} + \frac{\eta}{8}.$$

The exchangeability identity proved inside [Theorem 3](#) makes the inner sum $\mathbb{E}[1/(D+1)]$ at every iteration. With the chosen η , the last two terms sum to $\sqrt{\log(m)/(2N)}$. Finally, let $\bar{P} = N^{-1} \sum_{t=1}^N P_t$. The convexity established in the same proof gives $\Pr_v[c \succ_v \bar{P}^D] \leq N^{-1} \sum_t \Pr_v[c \succ_v P_t^D]$ for every c .

It remains to account for the running time. For a voter v and candidate c , let $b_{v,c}$ be the P_t -mass strictly below c . Because every $P_t(c)$ is positive, [Lemma 1](#) gives

$$g_{t,c} = \frac{1}{n} \sum_{v \in V} \sum_{d \in \text{supp}(D)} \Pr[D = d] \frac{(b_{v,c} + P_t(c))^{d+1} - b_{v,c}^{d+1}}{(d+1)P_t(c)}.$$

For each voter, all values $b_{v,c}$ are obtained in one scan of her ranking. Repeated squaring then evaluates the displayed powers in $O(nms \log(d_{\max} + 1))$ arithmetic operations per iteration; the multiplicative-weights update costs another $O(m)$ operations and exponential evaluations. This proves the claimed total. $\square$

For an infinite-support distribution, the same iteration guarantee holds provided that the required expectations can be evaluated to sufficient precision, but the running time then depends on the cost of those evaluations.

For a deterministic nonempty committee $A \subseteq C$, the event $A \succ_v (\text{RSL}_D)^D$ means that voter v 's favorite labeled member of A beats her favorite of the D independent draws.

Proposition 2. *Every nonempty committee $A \subseteq C$ satisfies*

$$\Pr_v[A \succ_v (\text{RSL}_D)^D] \leq \min \left\{ 1, |A| \mathbb{E} \left[\frac{1}{D+1} \right] \right\}.$$

Moreover,

$$\min_{P \in \Delta(C)} \max_{\emptyset \neq A \subseteq C} \frac{\Pr_v[A \succ_v P^D]}{|A|} = \mathbb{E} \left[\frac{1}{D+1} \right].$$

Proof. The event that A defeats $(\text{RSL}_D)^D$ is the union, over $c \in A$, of the event that the labeled copy of c defeats $(\text{RSL}_D)^D$. The union bound and Equation (3.1) give the first inequality. For any lottery P , the exchangeability identity in the proof of Theorem 3 gives some candidate c with $\Pr_v[c \succ_v P^D] \geq \mathbb{E}[1/(D+1)]$, which proves the matching lower bound for the normalized maximum. $\square$

Fix a committee A , expose the voter and the labeled copies in A , and let $T \in [0, 1]$ be the probability that one draw from RSL_D lies below the voter's favorite labeled member of A . Then $\Pr_v[A \succ_v (\text{RSL}_D)^r] = \mathbb{E}[T^r]$ and, by the conditioning calculation in Lemma 1, $\Pr_v[A \succ_v (\text{RSL}_D)^D] = \mathbb{E}[\phi(T)]$.

Because D is positive-integer-valued, ϕ is continuous and strictly increasing from 0 to 1 on $[0, 1]$, so its inverse $\phi^{-1} : [0, 1] \rightarrow [0, 1]$ is well-defined.

Lemma 7. *Let $s_A = \min \{1, |A| \mathbb{E}[1/(D+1)]\}$. For an integer $r \geq 1$, define $h_{r,\phi}(u) = (\phi^{-1}(u))^r$, and let $\text{cav } h_{r,\phi}$ be the smallest concave function on $[0, 1]$ lying above $h_{r,\phi}$. Then*

$$\Pr_v[A \succ_v (\text{RSL}_D)^r] \leq (\text{cav } h_{r,\phi})(s_A).$$

If $D = d$ deterministically, this bound is $s_A^{r/d}$ for $r \leq d$ and s_A for $r \geq d$.

Proof. By Proposition 2, $\mathbb{E}[\phi(T)] \leq s_A$. The function $\text{cav } h_{r,\phi}$ is nondecreasing, since it is at most the constant concave upper bound 1 and equals 1 at the right endpoint. Because $T^r = h_{r,\phi}(\phi(T))$, the definition of the concave upper bound and Jensen's inequality give

$$\mathbb{E}[T^r] \leq \mathbb{E}[(\text{cav } h_{r,\phi})(\phi(T))] \leq (\text{cav } h_{r,\phi})(\mathbb{E}[\phi(T)]) \leq (\text{cav } h_{r,\phi})(s_A).$$

For $\phi(t) = t^d$, the function $u^{r/d}$ is concave when $r \leq d$; when $r \geq d$, its smallest concave upper bound on $[0, 1]$ is the line u . $\square$

Let B be a lottery, let R be an independent positive integer, and write $\psi(t) = \mathbb{E}[t^R]$.

Theorem 5. *Every random-size stable lottery satisfies*

$$\Pr_v[B^R \succ_v (\text{RSL}_D)^D] \leq 1 - \psi \left(1 - \mathbb{E} \left[\frac{1}{D+1} \right] \right).$$

For deterministic $R = r$, the bound becomes $1 - (1 - \mathbb{E}[1/(D+1)])^r$.

Finally, random-size stable lotteries can be approximated in support, although the resulting additive error is too coarse for the main distortion proof. Let $\bar{d} = \mathbb{E}[D] = \phi'(1)$.

Proposition 3. *For every integer $s \geq 1$, there is a lottery $\hat{P}$ supported on at most s candidates such that*

$$\Pr_v[c \succ_v \hat{P}^D] \leq \mathbb{E} \left[\frac{1}{D+1} \right] + \bar{d} \sqrt{\frac{\pi}{2s}} \quad \text{for every } c \in C.$$

We now prove the last two statements and then record an extension to signed coefficients. The signed extension is mathematically useful but does not describe a random value of D unless the coefficients are nonnegative and sum to one.

B.3 A calculation with cumulative distributions

Draw a uniformly random voter v and order labeled candidates from worst to best according to v . For lotteries A and P , let $F_{A,v}$ and $F_{P,v}$ be their cumulative distribution functions in this order. Draw a labeled outcome ω from A and put

$$U = F_{P,v}(\omega), \quad Y = F_{A,v}(\omega).$$

Conditional on v , Y is uniform on $[0, 1]$. For every integer $d \geq 1$,

$$\Pr_v[A \succ_v P^d] = \mathbb{E}[U^d], \quad \Pr_v[A^d \succ_v P] = d \mathbb{E}[UY^{d-1}]. \quad (\text{B.1})$$

The first identity conditions on the draw from A . For the second, the best of d independent A -draws has density $dF_{A,v}^{d-1} dF_{A,v}$. Summing [Equation \(B.1\)](#) over the distribution of D , as in the conditioning calculation of [Lemma 1](#), gives the formulas used in [Lemma 8](#):

$$\Pr_v[A \succ_v P^D] = \mathbb{E}[\phi(U)], \quad \Pr_v[A^D \succ_v P] = \mathbb{E}[U\phi'(Y)].$$

Lemma 8. *For finitely supported D , define*

$$\mathcal{D}_\phi(u, v) = \phi(u) - \phi(v) - \phi'(v)(u - v).$$

Then

$$\Pr_v[A^D \succ_v P] = \Pr_v[A \succ_v P^D] + 1 - 2 \mathbb{E} \left[\frac{1}{D+1} \right] - \mathbb{E}[\mathcal{D}_\phi(U, Y)].$$

Consequently, $\Pr_v[A \succ_v P^D] \leq \mathbb{E}[1/(D+1)]$ implies $\Pr_v[A^D \succ_v P] \leq 1 - \mathbb{E}[1/(D+1)]$. The probability inequality extends to arbitrary D by truncation.

Proof. Convexity gives $\mathcal{D}_\phi \geq 0$, and rearranging its definition gives

$$U\phi'(Y) = \phi(U) - \mathcal{D}_\phi(U, Y) + Y\phi'(Y) - \phi(Y).$$

Because Y is uniform conditional on every voter,

$$\mathbb{E}[Y\phi'(Y) - \phi(Y)] = \int_0^1 (t\phi'(t) - \phi(t)) dt = 1 - 2 \int_0^1 \phi(t) dt = 1 - 2\mathbb{E}\left[\frac{1}{D+1}\right].$$

Substitution into the two preceding cumulative-distribution formulas proves the identity, and discarding the nonnegative term proves the consequence. Equality holds when $A = P$ by exchangeability. $\square$

B.4 Several draws from another lottery

For this proof, abbreviate $\mu = \mathbb{E}[1/(D+1)]$. Draw a labeled ω from B . The best of R draws from B has density $\psi'(F_{B,v}(\omega))$ with respect to a draw from B . Conditional on this best draw, the calculation in [Lemma 1](#) shows that all D draws from RSL_D lie below it with probability $\phi(F_{\text{RSL}_D,v}(\omega))$. With $U = F_{\text{RSL}_D,v}(\omega)$ and $Y = F_{B,v}(\omega)$, therefore,

$$\Pr_v[B^R \succ_v (\text{RSL}_D)^D] = \mathbb{E}[\phi(U)\psi'(Y)]. \quad (\text{B.2})$$

By linearity of [Equation \(3.1\)](#) in the candidate drawn from B ,

$$\mathbb{E}[\phi(U)] = \Pr_v[B \succ_v (\text{RSL}_D)^D] \leq \mu.$$

Also $0 \leq \phi(U) \leq 1$, and Y is uniform on $[0, 1]$. The function ψ' is nondecreasing. We claim that for every random pair (Z, Y) with $0 \leq Z \leq 1$, $\mathbb{E} Z \leq \mu$, and uniform marginal Y ,

$$\mathbb{E}[Z\psi'(Y)] \leq \int_{1-\mu}^1 \psi'(t) dt. \quad (\text{B.3})$$

To prove this directly, put $\vartheta = \psi'(1 - \mu)$. Pointwise,

$$Z(\psi'(Y) - \vartheta) \leq (\psi'(Y) - \vartheta)_+.$$

Taking expectations and using $\mathbb{E} Z \leq \mu$ gives

$$\mathbb{E}[Z\psi'(Y)] \leq \vartheta\mu + \int_0^1 (\psi'(t) - \vartheta)_+ dt = \int_{1-\mu}^1 \psi'(t) dt,$$

where a flat portion at the threshold does not affect the equality. Applying [Equation \(B.3\)](#) to $Z = \phi(U)$ in [Equation \(B.2\)](#) yields

$$\Pr_v[B^R \succ_v (\text{RSL}_D)^D] \leq \int_{1-\mu}^1 \psi'(t) dt = 1 - \psi(1 - \mu).$$

For deterministic $R = r$, substitute $\psi(t) = t^r$.

B.5 Approximating by a lottery with small support

Draw s candidates independently from RSL_D , and let $\hat{P}$ be their empirical distribution. Fix a voter. Place the candidates in this voter's order and linearly interpolate across the interval corresponding to each candidate; this interpolation is exactly the proportional label convention. Let F and $\hat{F}$ be the resulting cumulative distribution functions for RSL_D and $\hat{P}$. The sharp Dvoretzky–Kiefer–Wolfowitz inequality [Mas90] gives

$$\Pr \left[\sup_t |\hat{F}(t) - F(t)| > \varepsilon \right] \leq 2e^{-2s\varepsilon^2}.$$

Integrating the tail bound gives

$$\mathbb{E} \left[\sup_t |\hat{F}(t) - F(t)| \right] \leq \int_0^\infty 2e^{-2s\varepsilon^2} d\varepsilon = \sqrt{\frac{\pi}{2s}}. \quad (\text{B.4})$$

The same sampled candidates are used for every voter; Equation (B.4) is applied separately and then averaged over voters.

Because $\phi'(t) \leq \phi'(1) = \bar{d}$ on $[0, 1]$, the function ϕ is $\bar{d}$ -Lipschitz. By Lemma 1, for any pure candidate c , the probability that c defeats the empirical committee differs from the corresponding probability under RSL_D by at most $\bar{d} \sup_t |\hat{F}(t) - F(t)|$ for that voter. Therefore

$$\mathbb{E}_{\hat{P}} \left[\max_{c \in C} (\Pr_v[c \succ_v \hat{P}^D] - \Pr_v[c \succ_v (\text{RSL}_D)^D]) \right] \leq \bar{d} \sqrt{\frac{\pi}{2s}}.$$

The maximum causes no factor depending on $|C|$: for each voter, the same uniform error in the cumulative distributions controls every candidate position. Since $\Pr_v[c \succ_v (\text{RSL}_D)^D] \leq \mathbb{E}[1/(D+1)]$ for every c , some realization of the sample satisfies

$$\Pr_v[c \succ_v \hat{P}^D] \leq \mathbb{E} \left[\frac{1}{D+1} \right] + \bar{d} \sqrt{\frac{\pi}{2s}}$$

simultaneously for all candidates. Its empirical distribution has support at most s .

B.6 An extension to signed coefficients

Theorem 6. Let $(a_d)_{d \geq 1}$ be an absolutely summable sequence of real numbers, define

$$f(t) = \sum_{d \geq 1} a_d t^d,$$

and let $p_f = \int_0^1 f(t) dt$. Every finite profile admits a lottery P such that

$$\sum_{d \geq 1} a_d \Pr_v[c \succ_v P^d] \leq \int_0^1 f(t) dt$$

for every candidate c . This right-hand side is best possible: for every lottery P , the maximum of the left-hand side over candidates is at least $\int_0^1 f(t) dt$.

Proof. Within this proof, write $G_c(P)$ for the left-hand side. Since $p_f = \sum_{d \geq 1} a_d / (d + 1)$, absolute summability gives uniform convergence of the defining series on $[0, 1]$, hence continuity of every G_c . The exchangeability argument in the proof of [Theorem 3](#) is linear in the coefficients and gives

$$\sum_c P(c) G_c(P) = p_f.$$

For a lottery P , let $M(P) = \arg \max_{c \in C} G_c(P)$, and let $B(P)$ be the set of lotteries supported on $M(P)$. Each $B(P)$ is nonempty, convex, and compact. Continuity of the functions G_c implies that the correspondence B has a closed graph. Kakutani's fixed-point theorem therefore gives a lottery P with $P \in B(P)$. Every candidate in the support of this lottery maximizes $G_c(P)$, so the displayed identity gives

$$\max_c G_c(P) = \sum_c P(c) G_c(P) = p_f.$$

For every lottery, the maximum of the numbers $G_c(P)$ is at least their P -weighted average p_f , which proves the converse. $\square$

Signed coefficients do not define probabilities and need not make the weighted comparison function convex in P , so the minimax proof of [Theorem 3](#) does not apply. They also do not imply the union bound or the consequence of [Lemma 8](#) used in [Proposition 2](#). We use nonnegative coefficients summing to one everywhere in the distortion results.

C Proof of the Lower Bound

This section proves [Theorem 2](#). An *idealized profile* assigns arbitrary real-valued probabilities to a finite collection of rankings. We first show that a nondegenerate draw distribution admits at most one random-size stable lottery on every profile. We then give two idealized profile constructions whose distortion can be read exactly from the offset characterization, transfer their strict inequalities to ordinary finite electorates, and finish with an exact rational calculation showing that one of three profiles is hard, no matter how the number of draws is distributed.

Throughout this section, fix a random variable D supported on the positive integers. Recall its probability-generating function ϕ from [Definition 2](#), and its integral Φ and the benchmark $\mu = \Phi(1) = \mathbb{E}[1/(D + 1)]$ from [Definition 5](#).

For a lottery P and a candidate c , let $b_{v,c}(P)$ be the P -mass that voter v ranks strictly below c . If U is uniform on $[0, 1]$, then [Lemma 1](#) gives the comparison score of c against P as

$$L_c(P) = \mathbb{E}_{v,U} [\phi(b_{v,c}(P) + UP(c))] = \Pr_v[c \succ_v P^D].$$

Thus P is random-size stable if and only if $L_c(P) \leq \mu$ for every candidate c . By the self-play identity in the proof of [Theorem 3](#), $\sum_c P(c) L_c(P) = \mu$ for every P , and every candidate in the support of a stable lottery has score exactly μ .

C.1 Uniqueness and two hard-profile constructions

When D is not deterministically one, the function ϕ is strictly convex. This removes the otherwise problematic freedom to select among several stable lotteries.

Lemma 9 (Uniqueness). *If $\Pr[D \geq 2] > 0$, every probability distribution over the rankings of a finite candidate set, and in particular every finite profile, admits at most one random-size stable lottery for D .*

Proof. Suppose P and Q are stable and put $R = (P + Q)/2$. For every candidate c , both $b_{v,c}(P)$ and $P(c)$ are affine in P , so convexity of ϕ gives $L_c(R) \leq (L_c(P) + L_c(Q))/2 \leq \mu$. Hence R is stable. If $c \in \text{supp}(R)$, support tightness gives $L_c(R) = \mu$, so both inequalities are equalities. Strict convexity therefore implies

$$b_{v,c}(P) + UP(c) = b_{v,c}(Q) + UQ(c)$$

almost surely over v and U . Since U has a continuous distribution, the slopes agree and $P(c) = Q(c)$. This holds for every candidate in $\text{supp}(P) \cup \text{supp}(Q) = \text{supp}(R)$, while both lotteries vanish elsewhere. Thus $P = Q$. $\square$

The two constructions are most transparent through the following ratios. Their denominators are positive: for $p \in (0, 1)$, strict monotonicity of ϕ gives $\Phi(p)/p = \int_0^1 \phi(pt) dt < \int_0^1 \phi(t) dt = \mu$.

Definition 11: Ratios realized by the two constructions

For $r \in [0, 1]$, define

$$R_\phi(r) = r + \frac{\mu - \phi(r)}{1 - \mu}.$$

For $0 \leq a \leq p < 1$, define

$$S_\phi(a, p) = a + \frac{p(\Phi(a) + \mu - \Phi(1 + a - p)) - p\mu}{p\mu - \Phi(p)}.$$

Lemma 10 (Boundary construction). *Assume $\Pr[D \geq 2] > 0$. Let $r \in (0, 1)$ be rational, let N be a positive integer such that rN is integral, and suppose $\phi(r) < \mu$. There is an idealized distribution over rankings of $N + 1$ candidates whose unique random-size stable lottery P has offset value $1 + 2R_\phi(r)$.*

Proof. Let P be uniform on N exchangeable candidates and assign probability zero to one additional candidate o . Put $z = (\mu - \phi(r))/(1 - \phi(r)) \in (0, 1)$. Choose a uniform random order of the supported candidates. With probability z , place o above all of them; with probability $1 - z$, place o so that exactly rN supported candidates are below it.

Candidate o has comparison score $z + (1 - z)\phi(r) = \mu$. Symmetry gives the same score to every supported candidate, and their P -weighted average is the self-play value μ . Hence all candidates have score at most μ , so P is stable; it is unique by [Lemma 9](#).

Give o offset zero and assign offset one to every supported candidate. When o is ranked first, the excess-cost expression in [Theorem 4](#) is one and $\Delta_v = 0$; in the other rankings, they are r and one, respectively. The corresponding offset value is

$$1 + 2 \frac{z + (1 - z)r}{1 - z} = 1 + 2 \left(r + \frac{\mu - \phi(r)}{1 - \mu} \right) = 1 + 2R_\phi(r).$$

$\square$

Lemma 11 (Internal construction). *Assume $\Pr[D \geq 2] > 0$. Let $0 \leq a < p < 1$ be rational, let N be a positive integer such that aN and pN are integral, and suppose $\Phi(a) + \mu - \Phi(1 + a - p) > p\mu$. There is an idealized distribution over rankings of N candidates whose unique random-size stable lottery P has offset value $1 + 2S_\phi(a, p)$.*

Proof. Partition the candidates into a set H of size pN and a set L of size $(1 - p)N$, and let P be uniform on all N candidates. Put

$$K = p\mu - \Phi(p), \quad M = \Phi(a) + \mu - \Phi(1 + a - p), \quad z = \frac{K}{M - \Phi(p)}.$$

The hypothesis gives $M > p\mu > \Phi(p)$, so $z \in (0, 1)$. With probability z , place aN members of H at the bottom of the ranking, all members of L next, and the remaining members of H at the top. With probability $1 - z$, place all of H at the bottom and all of L at the top. In both cases, average uniformly over all permutations within each displayed group.

For any ranking, multiplying each candidate's score by its probability and summing over a set of candidates integrates ϕ over the percentile intervals occupied by that set. The integral over the positions occupied by H is M in the first ranking type and $\Phi(p)$ in the second. Its expectation is $zM + (1 - z)\Phi(p) = p\mu$. Symmetry therefore gives score μ to every member of H . The self-play identity then gives total weighted score $(1 - p)\mu$ to L , so symmetry gives score μ to every member of L . Thus P is stable, and it is unique by [Lemma 9](#).

Give candidates in H offset one and candidates in L offset zero, and designate any member of L as o . The first ranking type has excess-cost expression a and $\Delta_v = 1$, while the second has excess-cost expression p and $\Delta_v = 0$. Hence the corresponding offset value is

$$1 + 2 \frac{za + (1 - z)p}{z} = 1 + 2 \left(a + \frac{p(M - p\mu)}{K} \right) = 1 + 2S_\phi(a, p).$$

□

The preceding constructions use real-valued probabilities on a finite collection of rankings. The next lemma transfers a strict inequality to an ordinary finite electorate and makes the selection quantifier explicit.

Lemma 12 (Finite-profile transfer). *Suppose one of the idealized profiles in [Lemmas 10](#) and [11](#) has offset value strictly greater than $1 + 2c$. Then an ordinary finite profile on the same candidates has the property that every random-size stable lottery for D has distortion strictly greater than $1 + 2c$.*

Proof. Approximate the finitely many ranking probabilities by rational numbers and realize each rational distribution by a finite electorate. If the conclusion failed along every sufficiently accurate approximation, choose from each such profile a stable lottery whose offset value is at most $1 + 2c$. The candidate simplex is compact, so a subsequence of these lotteries converges. Candidate scores are continuous jointly in the ranking frequencies and the lottery, and therefore the limit is stable for the idealized profile. By [Lemma 9](#), it must be the uniform lottery displayed in the construction. The offset value is also continuous, and its denominator is positive at the idealized profile. The strict inequality must therefore hold throughout a sufficiently small neighborhood, a contradiction. Finally, [Theorem 4](#) turns this value into a lower bound on distortion. □

Lemma 13 (Candidate count). *For $r = 71/125$ and for $(a, p) \in \{(1/100, 37/50), (1/50, 3/4)\}$, the corresponding finite witness obtained from [Lemma 12](#) can be chosen with 5001 candidates.*

Proof. For every displayed parameter, $N = 5000$ is valid. The boundary construction then has 5001 candidates. The internal construction has 5000 candidates; add one candidate of probability zero at the bottom of every ranking and give it offset one. Its comparison score is $\phi(0) = 0$, and it changes neither stability nor the distortion calculation. Thus the finite transfer preserves a candidate set of size 5001. □

The case when $D = 1$ deterministically does not require the preceding uniqueness argument and admits a simple exact witness.

Lemma 14. *Every rule that returns a random-size stable lottery for deterministic $D = 1$ has distortion greater than $267/125$ on a profile with 5001 candidates.*

Proof. Place candidates a and b at 0 and 1 on the real line, respectively, and place the other 4999 candidates at 3. Of 1000 voters, place 501 at $499/1000$ and 499 at 1, breaking the metric ties among the additional candidates arbitrarily. Candidate a is the strict majority winner against b and is unanimously preferred to every additional candidate. For any lottery $P \neq \delta_a$, the comparison score of a against one draw from P is

$$\frac{1}{2}P(a) + \frac{501}{1000}P(b) + P(C \setminus \{a, b\}) > \frac{1}{2},$$

whereas δ_a is stable. Hence the unique stable lottery for one draw, equivalently the unique maximal lottery, selects a . Candidate b is socially optimal, while

$$SC(a) = \frac{748999}{10^6}, \quad SC(b) = \frac{251001}{10^6}, \quad \frac{SC(a)}{SC(b)} = \frac{748999}{251001} > \frac{267}{125}.$$

□

C.2 Why at least one of the three profiles exceeds the target

It remains to prove that, for every draw distribution, at least one of the three instances of the preceding constructions yields a ratio greater than $71/125$. The following proposition is the only computer-assisted step in the lower bound.

Proposition 4 (At least one of the three ratios exceeds the target). *For every positive-integer-valued random variable D , at least one of the following inequalities holds:*

$$\begin{aligned} R_\phi \left(\frac{71}{125} \right) &> \frac{71}{125}, \\ S_\phi \left(\frac{1}{100}, \frac{37}{50} \right) &> \frac{71}{125}, \\ S_\phi \left(\frac{1}{50}, \frac{3}{4} \right) &> \frac{71}{125}. \end{aligned}$$

Proof. Put $c = 71/125$ and let $v_d = \Pr[D = d]$. Define another probability distribution on the positive integers by

$$\hat{v}(d) = \frac{v_d}{\mu(d+1)}.$$

Indeed, its masses sum to one, and $\mathbb{E}_{d \sim \hat{v}}[d] = (1 - \mu)/\mu$ is finite and positive. For every positive integer d , define

$$b_d = (d+1)c^d - 1$$

and, for $0 \leq a \leq p < 1$, define

$$i_d(a, p) = p \left((c - a)(1 - p^d) - a^{d+1} - 1 + (1 + a - p)^{d+1} + p \right).$$

Substituting $\phi(x) = \sum_d \nu_d x^d$ and $\Phi(x) = \sum_d \nu_d x^{d+1} / (d+1)$ gives the exact identities

$$\mathbb{E}_{d \sim \hat{\nu}}[b_d] = \mathbb{E}_{d \sim \hat{\nu}}[d](c - R_\phi(c)), \quad (\text{C.1})$$

$$\mathbb{E}_{d \sim \hat{\nu}}[i_d(a, p)] = \mathbb{E}_{d \sim \hat{\nu}}[p(1 - p^d)](c - S_\phi(a, p)). \quad (\text{C.2})$$

Both multiplicative expectations on the right are strictly positive.

Consider the positive weights

$$(w_0, w_1, w_2) = \left(\frac{45697}{250000}, \frac{24337}{500000}, \frac{384269}{500000} \right),$$

which sum to one, and put $(a_1, p_1) = (1/100, 37/50)$ and $(a_2, p_2) = (1/50, 3/4)$. Exact rational arithmetic proves

$$F_d := w_0 b_d + w_1 i_d(a_1, p_1) + w_2 i_d(a_2, p_2) < 0 \quad \text{for every integer } d \geq 1. \quad (\text{C.3})$$

For completeness, we describe the finite verification and the infinite tail. Direct rational evaluation proves Equation (C.3) for $1 \leq d \leq 17$. Both parameter pairs satisfy $1 + a_j - p_j = 27/100$. Expanding the left-hand side gives

$$F_d = L + w_0(d+1)c^d + (w_1 p_1 + w_2 p_2) \left(\frac{27}{100} \right)^{d+1} - \sum_{j=1}^2 w_j \left((c - a_j) p_j^{d+1} + p_j a_j^{d+1} \right),$$

where $L = -w_0 + \sum_{j=1}^2 w_j p_j (c - a_j + p_j - 1)$. After discarding the nonpositive final sum, the two remaining positive exponential terms decrease for $d \geq 17$; for the first one, the ratio of successive terms is at most $(71/125)(19/18) < 1$. Exact rational evaluation at $d = 17$ makes the resulting upper bound smaller than $-1/15000$, proving Equation (C.3) for the entire tail.

The accompanying verifier `certificate_lower_bound.py` reconstructs the three functions from their definitions, performs every finite rational comparison, checks the tail expansion, and independently checks a second choice of positive weights that proves the same conclusion. The computation covers every possible positive number of draws, not merely a sampled or bounded set of values.

Because the three functions are bounded, averaging Equation (C.3) under $\hat{\nu}$ remains strictly negative even when D has infinite support. Consequently, at least one of the three expectations on the left of Equations (C.1) and (C.2) is negative. The positive factors on the right then imply that at least one of $R_\phi(c)$, $S_\phi(a_1, p_1)$, and $S_\phi(a_2, p_2)$ is strictly greater than c , as claimed. $\square$

Proof of Theorem 2. Fix the draw distribution of D before the profile is observed. If $D = 1$ deterministically, Lemma 14 gives the result. Otherwise, Proposition 4 makes one of the three displayed ratios strictly greater than $71/125$. In the boundary case, this inequality implies $\phi(71/125) < \mu$, so Lemma 10 applies. In either internal case, the first argument is smaller than $71/125$ and $p\mu - \Phi(p) > 0$, so the inequality implies $\Phi(a) + \mu - \Phi(1 + a - p) > p\mu$ and Lemma 11 applies. The resulting idealized profile has a unique stable lottery and offset value strictly greater than $1 + 2 \cdot 71/125 = 267/125$. By Lemmas 12 and 13, an ordinary finite profile with 5001 candidates has the same strict property for every random-size stable lottery. This proves the strict lower bound for each fixed draw distribution. Taking the infimum over all draw distributions turns these pointwise strict inequalities into the weak class lower bound $267/125$. $\square$