FOUR DUAL TUPLES SUFFICE: A SMALLER EXACT CERTIFICATE FOR THE RANDOMIZED METRIC-DISTORTION BOUND 11641/5000

KOREA SUPERINTELLIGENCE LABS

Abstract. Shah's recent bound of $11641/5000 = 2.3282$ on randomized metric distortion rests on an exact rational certificate: forty-nine parametric dual tuples for a semi-infinite linear program, certified feasible by 107 polynomial inequalities, and fifty candidate bounds whose disjunction is certified to cover the unit square by Bernstein-basis subdivision over 1096 accepted boxes. We show that the certificate is far smaller than it looks. *Four* of the fifty candidate bounds already cover the square — the tuples $\mathbf{z}_1, \mathbf{z}_2, \mathbf{z}_3$ and the tangency tuple at $\tau = 3/5$, with $\mathbf{z}_2$ used only where $x \geq 1/4$ — and the $(1-p)x \leq T(q)$ bound is not needed at all. The feasibility burden drops from 107 polynomial inequalities to 17. The four-tuple set is minimal for the subdivision search: dropping any one of the four makes it fail, and restoring the discarded bound does not repair the loss. Both halves of the reduced certificate are machine-checked in Lean 4. Two further facts about the published certificate are recorded: every one of its 107 feasibility polynomials passes the Bernstein test at the root, with no subdivision at all, so subdivision is needed only for the cover; and its own accepted cover invokes only 8 of its 50 candidates. We are careful about what is and is not verified: the passage from these polynomial inequalities to the distortion bound itself runs through four measure-theoretic steps of the source that are not formalized, and we say exactly which.

1. Introduction

1.1. **The objects.** In the metric model of social choice of Anshelevich, Bhardwaj, Elkind, Postl and Skowron [9], a finite set V of voters and a finite set C of m candidates are points of one unknown metric space, a voter's cost for a candidate is her distance to it, and each voter reports only the ranking of C that her distances induce. For a metric d consistent with the reported profile σ, the social cost of a candidate is $\mathrm{SC}_d(c) = \mathbb{E}_v[d(v,c)]$, and that of a lottery P over candidates is $\mathrm{SC}_d(P) = \mathbb{E}_{c\sim P}\,\mathrm{SC}_d(c)$. The *distortion* of P at σ is

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

and the distortion of a randomized voting rule is the supremum of this over profiles. The optimal deterministic distortion is 3 [8]. For randomized rules the best known lower bound is $2.1126\ldots$ [10], and the upper bound went 2.753 [4] $\to$ 5/2 [5, 6] $\to$ $11641/5000 = 2.3282$, the last in Shah [1], the source of this note.

The 5/2 bound and the bound of [1] are both proved through the candidate-set reformulation of Charikar, Ramakrishnan, Wang and Wu [4], in the form made explicit by Frank [5]: for a nonempty proper $I \subsetneq C$ with $J = C \setminus I$, set

$$q(I) = \min_{i\in I} \Pr_v[\, j \succ_v i \text{ for some } j \in J\,], \qquad \ell(P,I) = \sum_{j\in J} P(j)\,\Pr_v[\, i \succ_v j \text{ for every } i \in I\,], \tag{1}$$

and then $\ell(P,I) \leq \kappa\, q(I)$ for every such I implies $\mathrm{dist}(P) \leq 1 + 2\kappa$. The new ingredient of [1] is the *random-size stable lottery* $\mathrm{RSL}(D)$: for a positive-integer-valued random variable D it is a lottery satisfying $\mathbb{E}_D \Pr_v[A \succ_v (\mathrm{RSL}(D))^D] \leq \mathbb{E}[1/(D+1)]$ for every challenger lottery A, where $(\mathrm{RSL}(D))^D$ is a multiset of D independent draws compared by random-label tie-breaking. Deterministic $D = 1$ recovers maximal lotteries and deterministic $D = k$ the stable k-lotteries of Charikar, Ramakrishnan, Tan and Wang [7]. With $D^\star$ supported on $\{1, 12, 13\}$ and the mixture $P^\star = \frac{7941}{50000}\mathrm{IV} + \frac{42059}{50000}\mathrm{RSL}(D^\star)$ of

Integrated Veto with RSL($D^\star$), the bound 2.3282 follows from a single scalar inequality, [1, Lemma 5]: $g_{D^\star}(q) \le T(q) = q(a - bq)$ for all $q \in [0, 1]$, where $a = 33205/42059$, $b = 7941/42059$ and g_D is a worst-case loss envelope defined by an optimization over a scalar and two Borel measures.

1.2. **The certificate, and the question this note answers.** Lemma 5 of [1] is proved by an exact rational certificate with two halves. Dualizing the envelope program at each fixed (q, x) produces a semi-infinite linear program Dual(D, q, x) with four scalar variables $\mathbf{z} = (\lambda, \mu, u, v)$; the source exhibits *forty-nine* parametric dual tuples $\mathbf{z}_j(x)$, and

- *dual feasibility*: each $\mathbf{z}_j$ is certified feasible on a stated range of x by a list of polynomial inequalities $H^{\text{feas}} \ge 0$ — 107 of them in total;
- *the cover*: the *fifty* candidate bounds obtained from these tuples together with the trivial bound $(1 - p)x \le T(q)$ have a disjunction that is certified to hold at every (q, x) of the unit square, by the Bernstein-basis subdivision test of Garloff [3] over 1096 accepted rational boxes, together with a second, separate verification on the line $x = 1$.

The verifier is exact-rational Python [2]; there was no formal artifact. Box counts quoted for [1], here and below, are not stated in that paper: they are measured from a run of its own verifier (§6). Fifty candidate bounds and 107 feasibility polynomials is a large object, and the source's own discussion [1, §6] lists four avenues for improving the constant without asking how much of the certificate is load-bearing. That is the question here: *how small can the certificate be made, and which of its parts do any work?*

1.3. **What is settled.** The answer is that almost none of it is needed.

- **Four candidate bounds cover the square** (Theorem 5.1). At every $(q, x) \in [0, 1]^2$ at least one of $H_1^{\text{out}}, H_2^{\text{out}}, H_3^{\text{out}}, H_{49}^{\text{out}}$ is nonnegative, with H_2^{out} — whose tuple is feasible only for x bounded away from 0 — used only where $x \ge 1/4$. The trivial $(1-p)x \le T(q)$ bound is never used. Only 17 of the 107 feasibility polynomials survive (Theorems 4.1–4.4), and the accepted subdivision has 1022 boxes against the published 1096.
- **Four is minimal for this search** (§6). Dropping any one of $\mathbf{z}_1, \mathbf{z}_2, \mathbf{z}_3, \mathbf{z}_{49}$ makes the subdivision search fail, and restoring the discarded $(1 - p)x \le T(q)$ bound does not repair the loss of $\mathbf{z}_{49}$. One instance of this is itself machine-checked (Proposition 7.1(iii)).
- **Two facts about the published certificate** (§6). Every one of its 107 dual-feasibility polynomials passes the Bernstein test at the root of the pipeline, with no subdivision at all: the subdivision loop is needed only for the cover. And the published accepted cover fires only 8 of its 50 candidates.
- **Everything above is kernel-checked**, in Lean 4 against *Mathlib*, with the certificate polynomials rebuilt inside the proof from the source's formulas rather than transcribed, and with every exact polynomial division that defines a H^{feas} or a H^{out} verified as an identity rather than assumed (§9).

The reduction changes no constant: the bound proved is the source's 11641/5000, and the independent check is the source's own verifier, which accepts the reduced certificate on a reduced input file with no edit to its code (§6).

1.4. **What is not claimed.** This note verifies the *polynomial* content of [1, Lemma 5] on the region where the source needs it. It does not verify dist$(P^\star) \le 11641/5000$. Four steps of the source stand between the two, and none of them is a computation: the set-certificate proposition, the measure-valued loss-envelope lemma, generalized-moment weak duality together with the sign chain that turns the $H^{\text{feas}} \ge 0$ inequalities into dual feasibility, and the existence of RSL($D^\star$). Section 8 names them one by one. The artifact here should be read as *the computer-assisted half of the source's proof, redone in a kernel with a certificate an order of magnitude smaller*, and not as a formal proof of 2.3282.

1.5. **Method.** Two ingredients make the reduction affordable. The first is a rewriting of the certificate's own algebra: the binomial double sums (B.1)–(B.3) of [1, App. B.3] all collapse to differences of $\Psi(s) = \int_0^s f$, so the entire certificate is generated by f, f' and Ψ with no binomial coefficient anywhere (Lemma 3.1; proved on paper, outside the formal development). The second is the shape of the Bernstein test itself: after homogenization, "all Bernstein coefficients are nonnegative" is literally "this polynomial in $2n$ variables has nonnegative coefficients", so the coefficients are the *output* of the same additions and multiplications the defining equations use, and never an input to be trusted. Two additions to that instrument were needed and are new here: nonnegativity on a dyadic *slab* of one axis, for the tuple $\mathbf{z}_2$, which is genuinely negative near $x = 0$; and a subdivision tree for a *disjunction*, whose leaves each name the member of the candidate family that certifies them and whose interior nodes prune members used nowhere below them.

2. Preliminaries

All of this section is the source's, in the source's notation; nothing here is new.

2.1. **Candidate sets.** With $q(I)$ and $\ell(P, I)$ as in (1), the reformulation used throughout is:

Proposition 2.1 ([4]; in this form [5])**.** *Let $\kappa \geq 0$. If a lottery P satisfies $\ell(P, I) \leq \kappa\, q(I)$ for every nonempty proper $I \subsetneq C$, then* $\operatorname{dist}(P) \leq 1 + 2\kappa$.

Frank [5] and Ye [6] reach $\kappa = 3/4$, hence $5/2$, with the equal mixture of a maximal lottery and Integrated Veto, which turns the Simultaneous Plurality Veto process of Kizilkaya and Kempe [11] into a lottery. Proposition 2.1 yields nothing better from any mixture of those two rules, and Ye [6] proves a stronger obstruction: a much larger family of rules has distortion no better than $5/2$ [1, §3]. The source's target is $\kappa = 6641/10000$.

2.2. **The loss envelope and its dual.** Fix D, write $f = \phi_D$ for its probability generating function, $p = \int_0^1 f$ and $r = 1 - x$. For $x \in (0, 1]$ put

$$K_x = \int_0^1 f(ru)\, du, \qquad W_x(y) = f(x + ry), \qquad w_x(b_1) = f(xb_1),$$

$$G_x(b_0, y) = \frac{\int_0^y f(ru + xb_0)\, du + \int_y^1 f(ru + x)\, du - K_x}{x}, \tag{2}$$

$$F_x(b_1) = \frac{\int_0^1 f(ru + xb_1)\, du - K_x}{x}, \qquad R_x = \frac{p - K_x}{x}.$$

The source's Lemma 2 bounds $\ell(\mathrm{RSL}(D), I)$ by $g_D(q(I))$, where $g_D(q)$ is the value of a program over a scalar $x \in [0, 1]$ and two Borel probability measures η on $[0, 1]^2$ and ξ on $[0, 1]$ subject to three constraints (C1)–(C3); its Lemma 3 normalizes the program at fixed x into a linear program $\mathrm{Primal}(D, q, x)$ in (η, ξ) with the two constraints built from (2); and its Lemma 4 exhibits the dual

$$\mathrm{Dual}(D, q, x): \qquad \inf_{\lambda, \mu \geq 0,\ u, v \in \mathbb{R}} \lambda p + \mu R_x + (1 - q)u + qv$$

subject to $u + \lambda W_x(y) + \mu G_x(b_0, y) \geq b_0$ for all $(b_0, y) \in [0, 1]^2$ and $v + \lambda w_x(b_1) + \mu F_x(b_1) \geq b_1$ for all $b_1 \in [0, 1]$, a semi-infinite linear program with four scalar variables. Weak duality for generalized moment problems [12] makes every feasible tuple a certificate: writing

$$U(q, x; \mathbf{z}) = x\big(\lambda p + \mu R_x + (1 - q)u + qv\big) \qquad \text{for } \mathbf{z} = (\lambda, \mu, u, v), \tag{3}$$

the source's Lemma 4 reduces $g_D(q) \leq T(q)$ to the following disjunction at each $q \in [0, 1]$ and $x \in (0, 1]$: either $\mathrm{Primal}(D, q, x)$ is infeasible, or $(1 - p)x \leq T(q)$, or some feasible dual tuple $\mathbf{z}$ has $U(q, x; \mathbf{z}) \leq T(q)$.

2.3. **The rule, the target and the four tuples.** The source takes $D^\star$ supported on $\{1,12,13\}$ with $\Pr[D^\star=1]=423239/10^6$, $\Pr[D^\star=12]=471105/10^6$, $\Pr[D^\star=13]=105656/10^6$, so that

$$f(t)=\phi_{D^\star}(t)=\frac{423239\,t+471105\,t^{12}+105656\,t^{13}}{10^6},\qquad p=\int_0^1 f=\frac{46483747}{182000000}, \tag{4}$$

and the target $T(q)=q(a-bq)$ with $a=33205/42059$, $b=7941/42059$. Of its 49 parametric tuples we shall need exactly four. Write $\Psi(s)=\int_0^s f$ and $L_x=p-\Psi(x)-\Psi(r)-xf(r)$.

- $\mathbf{z}_1$ and $\mathbf{z}_2$ (*dual families 1 and 2*), with $y_1=80529/100000$, $y_2=1$, $A_x=G_x(1,0)$ and
$$\Delta_j(x)=\big(f(x)-W_x(y_j)\big)\partial_y G_x(0,y_j)-\big(A_x-G_x(0,y_j)\big)W_x'(y_j),$$
by $\lambda_j=\partial_y G_x(0,y_j)/\Delta_j(x)$, $\mu_j=-W_x'(y_j)/\Delta_j(x)$, $u_j=-\lambda_j W_x(y_j)-\mu_j G_x(0,y_j)$, $v_j=0$. Their stated eligibility ranges are $x\in(0,1)$ for $\mathbf{z}_1$ and $x\in[1/5,1)$ for $\mathbf{z}_2$.
- $\mathbf{z}_3$ (*dual family 3*), by $\lambda_3=L_x/\big(xf(r)(1-f(x))\big)$, $\mu_3=r/f(r)$, $u_3=-\lambda_3$, $v_3=0$, eligible on $x\in(0,1)$.
- $\mathbf{z}_{49}$, the tangency tuple at $\tau_{49}=(49+11)/100=3/5$: with $S_j(x)=\lambda_3 xf'(x\tau_j)+\mu_3F_x'(\tau_j)$ and $\theta_j=1/S_j$, $\lambda_j=\theta_j\lambda_3$, $\mu_j=\theta_j\mu_3$, $u_j=(1-\theta_j)+\theta_j u_3$, $v_j=\tau_j-\lambda_j f(x\tau_j)-\mu_j F_x(\tau_j)$, eligible on $x\in(0,1]$.

The polynomials attached to these tuples are the source's. For $j\in\{1,2\}$, with $I_{j,0}=u_j+\lambda_j W_x(y)+\mu_j G_x(0,y)$ and $I_{j,1}=I_{j,0}+\mu_j f(ry)/r-1$, they are

$$H_{1,j}^{\mathrm{feas}}=-\frac{\Delta_j(x)}{x(1-x)^2},\qquad H_{2,j}^{\mathrm{feas}}=-\partial_y G_x(0,y_j),\qquad H_{3,j}^{\mathrm{feas}}=W_x'(y_j),$$
$$H_{4,j}^{\mathrm{feas}}=\mathrm{num}\Big(\lambda_j a_1 x+\mu_j\frac{f(r)}{r}-1\Big),\qquad H_{5,j}^{\mathrm{feas}}=\frac{\mathrm{num}(I_{j,0})}{(y-y_j)^2},\qquad H_{6,j}^{\mathrm{feas}}=\frac{\mathrm{num}(I_{j,1})}{y}, \tag{5}$$

with $a_1=423239/10^6$ and num the numerator in lowest terms, signed so that the denominator is positive on the domain. For $\mathbf{z}_3$, with $I_{3,0}=\mu_3G_x(0,y)-\lambda_3(1-W_x(y))$ and $I_{3,1}=I_{3,0}+\mu_3 f(ry)/r-1$,

$$H_{1,3}^{\mathrm{feas}}=\frac{L_x}{x(1-x)},\qquad H_{2,3}^{\mathrm{feas}}=\frac{\mathrm{num}(I_{3,0})}{1-y},\qquad H_{3,3}^{\mathrm{feas}}=\frac{\mathrm{num}(I_{3,1})}{y(1-y)}; \tag{6}$$

and for $\mathbf{z}_{49}$, with $Z_j(x)=L_xf'(x\tau_j)+(1-f(x))\big(f(r+x\tau_j)-f(x\tau_j)\big)$,

$$H_{1,j}^{\mathrm{feas}}=\frac{Z_j(x)}{1-x},\qquad H_{2,j}^{\mathrm{feas}}=\frac{Z_j(x)-(1-f(x))f(1-x)}{x(1-x)}. \tag{7}$$

Finally the *outer* polynomials, one per candidate bound, are the cleared forms of $T(q)-U(q,x;\mathbf{z}_j)$:

$$\begin{aligned}H_j^{\mathrm{out}}&=-\frac{\Delta_j(x)}{x}\big(T(q)-U(q,x;\mathbf{z}_j)\big)\quad(j\in\{1,2\}),\\ H_3^{\mathrm{out}}&=f(r)\big(1-f(x)\big)\big(T-U\big),\qquad H_{49}^{\mathrm{out}}=\frac{Z_{49}(x)}{1-x}\big(T-U\big),\end{aligned} \tag{8}$$

each of which is a polynomial in (q,x) with rational coefficients; the displayed prefactors are positive on the relative interior of the region where the corresponding tuple is used, so $H_j^{\mathrm{out}}\ge 0$ is exactly the candidate bound $U(q,x;\mathbf{z}_j)\le T(q)$ there.

2.4. **Normalization.** Every polynomial below is understood in a fixed *integer-cleared* normalization: it is multiplied by an explicit positive integer — powers of 10^6 and 364, and powers of the denominators 10^5 of y_1 and 5 of τ_{49} — so that all its coefficients are integers. Multiplying by a positive constant changes no sign claim. The normalization is not left to convention: each theorem below states the exact divisor of the division that defines the polynomial, and asserts the corresponding identity against a numerator assembled from the formulas of §2, so the constant is pinned by the identity. In the same normalization the source's strict claims $H_{1,j}^{\mathrm{feas}}>0$ become the quantitative $H_{1,j}^{\mathrm{feas}}\ge 1$.

One discrepancy inside [1] decides which polynomial the theorems of §4 are about, so we record it. The convention just quoted asks for num *in lowest terms*, but the published verifier [2] does not reduce: for six of the seventeen polynomials the cleared numerator it builds carries an extra factor,

$x(1-x)^2$ for $H^{\mathrm{feas}}_{5,1}, H^{\mathrm{feas}}_{5,2}, H^{\mathrm{feas}}_{2,3}, H^{\mathrm{feas}}_{3,3}$ and $1-x$ for $H^{\mathrm{feas}}_{6,1}, H^{\mathrm{feas}}_{6,2}$. Both are nonnegative on $[0,1]$ and positive on $(0,1)$, so the two forms have the same sign wherever the tuples are used, and the sign chain of [1, App. B.3], which needs $I_{j,0}, I_{j,1} \geq 0$ there, is served by either. We follow the verifier, so that the polynomials proved nonnegative below are literally the ones the published run tests; the other eleven agree with the prose convention up to a positive integer.

3. The generating identities

Appendix B.3 of [1] rewrites the kernels (2) as binomial double sums: for $f(t) = \sum_d a_d t^d$,

$$F_x(b_1) = \sum_d a_d \sum_{k=1}^{d} \binom{d}{k} \frac{x^{k-1} r^{d-k}}{d-k+1} b_1^k,$$

and similarly for $G_x(b_0, y)$ and R_x. These sums are what the published verifier manipulates, with the one exception noted at the end of this section. They are unnecessary: every one of them collapses to a difference of values of the single antiderivative $\Psi(s) = \int_0^s f$.

Lemma 3.1. *Let f be any polynomial, $p = \int_0^1 f$, $\Psi(s) = \int_0^s f$ and $r = 1-x$. Then for $x \in (0,1)$ and all $y, b \in [0,1]$,*

$$\begin{aligned}
rx\, G_x(0,y) &= p - \Psi(r) + \Psi(ry) - \Psi(x+ry), & x\, \partial_y G_x(0,y) &= f(ry) - f(x+ry),\\
rx\, G_x(1,0) &= p - \Psi(x) - \Psi(r), & rx\, F_x(b) &= \Psi(r+xb) - \Psi(xb) - \Psi(r),\\
rx\, R_x &= rp - \Psi(r), & W_x(y) &= f(x+ry), \quad W_x'(y) = r\, f'(x+ry).
\end{aligned}$$

Proof. Substituting $t = ru$ gives $K_x = \int_0^1 f(ru)\, du = \Psi(r)/r$, and likewise $\int_0^y f(ru)\, du = \Psi(ry)/r$ and $\int_y^1 f(ru+x)\, du = \big(\Psi(r+x) - \Psi(ry+x)\big)/r = \big(p - \Psi(x+ry)\big)/r$, since $r + x = 1$. Putting these into the definition of $G_x(0,y)$ in (2) and multiplying by rx gives the first identity; differentiating the same numerator in y under the integral sign gives the second. Taking $b_0 = 1$ and $y = 0$ leaves $xG_x(1,0) = \int_0^1 f(ru+x)\, du - K_x = \big(p - \Psi(x)\big)/r - \Psi(r)/r$, which is the third. The fourth is $xF_x(b) = \int_0^1 f(ru+xb)\, du - K_x = \big(\Psi(r+xb) - \Psi(xb)\big)/r - \Psi(r)/r$, and the fifth is $xR_x = p - K_x = p - \Psi(r)/r$. The last two are the definitions. □

Lemma 3.1 is proved here on paper and is *not* part of the formal development. Every theorem below is an unconditional statement about explicit polynomials, and none of them depends on it; what does depend on it is the *reading* of those polynomials as the certificate objects of [1, App. B.3], since the formal development builds them in the Ψ form rather than the binomial form. Its role is otherwise practical, and decisive. Every object of the certificate is a short expression in f, f' and Ψ — three explicit polynomials of degree at most 14 for the f of (4) — with one explicit integer clearing constant. No binomial coefficient and no division by $d-k+1$ appears anywhere, and the polynomials of (5)–(8) can therefore be *built* inside the proof by a handful of additions and multiplications, instead of being transcribed from the verifier's output. That is what makes a kernel-level check of the certificate affordable at all; without it the same objects arrive as opaque coefficient arrays.

The identities are one substitution each and nothing about them is claimed to be new; the point is that the source does not use them. In the text of [1] the antiderivative Ψ occurs in exactly one place, the definition of $L_x = p - \Psi(x) - \Psi(r) - xf(r)$ for $\mathbf{z}_3$, and the kernels (2) are stated and used as the binomial sums (B.1)–(B.3). The verifier [2] has one further occurrence — its `dual_family_3` builds the G-numerator in the shape of the first identity above — and works from (B.1)–(B.3) everywhere else.

Each identity of Lemma 3.1 was in addition checked numerically against the published verifier's own binomial-sum objects at three rational points before any of it was used, and each of the 21 numerators rebuilt from f, f', Ψ below was checked to agree with the verifier's corresponding polynomial coefficient by coefficient, with residual exactly zero.

4. Dual feasibility of the four tuples

We now state what is proved. In each theorem the H^{feas} are the polynomials (5)–(7) in the normalization of §2.4, and N_i denotes the numerator built from the formulas of §2 through Lemma 3.1. The second half of each conclusion — the exact-division identity — is what makes the first half a statement about the source's certificate rather than about a transcription of it: a wrong coefficient in the quotient breaks the identity instead of weakening the inequality.

Theorem 4.1 (dual family 1, $y_1 = 80529/100000$)**.** *For all $(x, y) \in [0,1]^2$ the six polynomials $H^{\mathrm{feas}}_{i,1}$, $1 \le i \le 6$, satisfy*

$$H^{\mathrm{feas}}_{1,1} \ge 1, \qquad H^{\mathrm{feas}}_{i,1} \ge 0 \quad (2 \le i \le 6),$$

and each is the exact quotient recorded by its defining division:

$$x^2(1-x)^3\, H^{\mathrm{feas}}_{1,1} = N_1, \quad x\, H^{\mathrm{feas}}_{2,1} = N_2, \quad H^{\mathrm{feas}}_{3,1} = N_3,$$

$$x(1-x)\, H^{\mathrm{feas}}_{4,1} = N_4, \quad x(1-x)(10^5 y - 80529)^2\, H^{\mathrm{feas}}_{5,1} = N_5, \quad x(1-x)y\, H^{\mathrm{feas}}_{6,1} = N_6.$$

Theorem 4.2 (dual family 2, $y_2 = 1$)**.** *For all (x, y) with $1/4 \le x \le 1$ and $0 \le y \le 1$ the six polynomials $H^{\mathrm{feas}}_{i,2}$ satisfy $H^{\mathrm{feas}}_{1,2} \ge 1$ and $H^{\mathrm{feas}}_{i,2} \ge 0$ for $2 \le i \le 6$, and*

$$x^2(1-x)^3\, H^{\mathrm{feas}}_{1,2} = N_1, \quad x\, H^{\mathrm{feas}}_{2,2} = N_2, \quad H^{\mathrm{feas}}_{3,2} = N_3,$$

$$x(1-x)\, H^{\mathrm{feas}}_{4,2} = N_4, \quad x(1-x)(y-1)^2\, H^{\mathrm{feas}}_{5,2} = N_5, \quad x(1-x)y\, H^{\mathrm{feas}}_{6,2} = N_6.$$

Theorem 4.3 (dual family 3)**.** *For all $(x, y) \in [0,1]^2$, $H^{\mathrm{feas}}_{1,3} \ge 0$, $H^{\mathrm{feas}}_{2,3} \ge 0$, $H^{\mathrm{feas}}_{3,3} \ge 0$, and $x(1-x)\, H^{\mathrm{feas}}_{1,3} = N_1$, $(1-y)\, H^{\mathrm{feas}}_{2,3} = N_2$, $y(1-y)\, H^{\mathrm{feas}}_{3,3} = N_3$.*

Theorem 4.4 (the tangency tuple at $\tau = 3/5$)**.** *For all $(x, y) \in [0,1]^2$, $H^{\mathrm{feas}}_{1,49} \ge 1$ and $H^{\mathrm{feas}}_{2,49} \ge 0$, and $(1-x)\, H^{\mathrm{feas}}_{1,49} = N_1$, $x(1-x)\, H^{\mathrm{feas}}_{2,49} = N_2$.*

These are 17 polynomial inequalities against the 107 of the published certificate; the 107 is $6+6+3$ for the first three families plus 2 for each of the 46 tangency tuples $\mathbf{z}_4, \dots, \mathbf{z}_{49}$.

Two remarks on the domains. Theorem 4.2 restricts $\mathbf{z}_2$ to $x \ge 1/4$ where the source proves feasibility on $x \ge 1/5$; since $[1/4, 1] \subset [1/5, 1]$, the statement is a *restriction* of the source's, not an extension. The threshold moved because $1/5$ is not dyadic and the subdivision instrument used here splits at midpoints; the slab $[1/4, 1]$ is reached from the unit box by two half-box substitutions. The restriction costs nothing: the cover of §5 still closes with $\mathbf{z}_2$ confined to $x \ge 1/4$. Neither $1/5$ nor $1/4$ is the true threshold — the same six polynomials were observed to pass the Bernstein test on $[1/8, 1/4]$ as well — but nothing below needs that. Second, the slab hypothesis is not vacuous: $H^{\mathrm{feas}}_{5,2}$ is strictly negative at $(x, y) = (1/16, 15/16)$ (Proposition 7.1(ii)), so the conclusion of Theorem 4.2 genuinely fails somewhere in the unit box.

Finally, the sign conclusions above are the sign conclusions of the published verifier — equivalently, on the relative interior, the source's, the two forms differing only by the factors recorded in §2.4 — and the passage from them to the statement "$\mathbf{z}_j$ is dual-feasible" is the source's prose argument from convexity, which is not formalized here; see §8.

5. The cover on four candidate bounds

Theorem 5.1 (the reduced cover)**.** *For every $(q, x) \in [0,1]^2$,*

$$H^{\mathrm{out}}_1(q,x) \ge 0 \quad \vee \quad \Big(x \ge \tfrac{1}{4} \ \wedge \ H^{\mathrm{out}}_2(q,x) \ge 0\Big) \quad \vee \quad H^{\mathrm{out}}_3(q,x) \ge 0 \quad \vee \quad H^{\mathrm{out}}_{49}(q,x) \ge 0.$$

Proposition 5.2 (the four outer divisions)**.** *Identically in $(q, x) \in \mathbb{R}^2$,*

$$x^2(1-x)\, H^{\mathrm{out}}_1 = N_1, \qquad x^2(1-x)\, H^{\mathrm{out}}_2 = N_2, \qquad H^{\mathrm{out}}_3 = N_3, \qquad (1-x)\, H^{\mathrm{out}}_{49} = N_{49},$$

where N_j is the numerator built from the source's formulas for T, for $U(\cdot; \mathbf{z}_j)$ of (3) and for the prefactor of (8).

The published cover uses fifty candidate bounds — the $(1-p)x \le T(q)$ bound and all 49 tuples — over 1096 accepted boxes. Theorem 5.1 uses four, over 1022, and never uses the $(1-p)x \le T(q)$ bound.

Proof sketch, and what rests on exhaustive computation. The whole of Theorem 5.1 rests on exhaustive computation, and the shape of that computation is the following. Write a polynomial H of bidegree (m, n) in the tensor-product Bernstein basis on a box; since the basis functions are nonnegative and sum to 1, nonnegative Bernstein coefficients prove $H \ge 0$ throughout the box. The device used here computes those coefficients without a basis conversion: with $u_j = x_j$ and $v_j = 1 - x_j$, the homogenization $\widehat{H}(u, v) = \prod_j (u_j + v_j)^{n_j} H(u_j/(u_j + v_j))$ has monomial coefficients equal to the Bernstein coefficients up to positive binomial factors, and $\widehat{H}(t, 1-t) = H(t)$; homogenization is multiplicative, so $\widehat{H}$ is assembled by the same additions and multiplications that define H, and "all Bernstein coefficients are nonnegative" becomes "this integer coefficient list has no negative entry". Subdivision is free in the same representation, because the two half-box substitutions $t_j \mapsto t_j/2$ and $t_j \mapsto (1 + t_j)/2$ are integer linear substitutions with nonnegative coefficients.

Theorem 5.1 is proved by splitting the square at $x = 1/4$ and $x = 1/2$ into three dyadic slabs and walking, on each, a binary subdivision tree whose leaves each name the candidate that certifies them and whose interior nodes carry the list of candidates still live below them. The finite search that was run, and its result, is exactly:

slab	candidates offered	accepted boxes	depth
$x \in [0, 1/4]$	$\mathbf{z}_1, \mathbf{z}_3, \mathbf{z}_{49}$	54	29
$x \in [1/4, 1/2]$	$\mathbf{z}_1, \mathbf{z}_2, \mathbf{z}_3, \mathbf{z}_{49}$	11	8
$x \in [1/2, 1]$	$\mathbf{z}_1, \mathbf{z}_2, \mathbf{z}_3, \mathbf{z}_{49}$	957	32

$\mathbf{z}_2$ is absent from the first slab, which is what puts the side condition $x \ge 1/4$ on the second disjunct of Theorem 5.1 and matches it to the domain of Theorem 4.2. Each of the 1022 leaves is one evaluation of the Boolean "this coefficient list has no negative entry"; the soundness of the tree walk — that an accepted tree forces the disjunction at every point of the slab, and that pruning a candidate at an interior node only weakens the claim — is proved once and for all, without any compiled evaluation. Theorems 4.1–4.4 rest on exhaustive computation in the same sense, but far more cheaply: each of the 17 polynomials passes the test at the *root*, with no subdivision, once the domain is fixed (for dual family 2, after the two half-box steps that cut the slab $[1/4, 1]$). The exact-division identities of Theorems 4.1–4.4 and Proposition 5.2 are decided by the same kind of computation: the term list of divisor × quotient − numerator is computed and compared with the empty list.

Corollary 5.3. *Restricted to $(q, x) \in [0, 1] \times (0, 1)$, the disjunction of Theorem 5.1 invokes each tuple only on a range of x on which* [1, Tab. 4] *states it feasible, and for each of those four tuples the polynomial signs that* [1, App. B.3] *asks for on that range are the conclusions of Theorems 4.1–4.4. On that region the pair (Theorems 4.1–4.4, Theorem 5.1) therefore carries the entire computational content of the proof of* [1, Lemma 5]*, with four candidate bounds and* 17 *feasibility polynomials in place of fifty and* 107*.*

The region $[0, 1] \times (0, 1)$ is the first of the two disjunctions the source verifies. The second, the line $x = 1$ with $q \in [1 - p, 1]$, is not reduced here: the source treats it with the $(1-p)x \le T(q)$ bound and the tangency tuples, and the accompanying primal-infeasibility argument at $x = 1$ is prose (§8). Theorem 5.1 does hold on the closed square, including $x \in \{0, 1\}$; we make no claim about tuple eligibility there.

6. The certificate, measured

The statements in this section are *measurements* of the published certificate and of the subdivision search, obtained by instrumenting the published verifier [2] from a separate script without editing it. They are not theorems, and none of them is machine-checked, with the one exception noted below.

The published run. Reproducing the published verification takes 5 min 54 s of wall time and 345 CPU-seconds in exact rational arithmetic. Recording which candidate certifies each of its 1096 accepted boxes gives the census

$$(\mathrm{C3}) : 2, \quad \mathbf{z}_1 : 44, \quad \mathbf{z}_2 : 513, \quad \mathbf{z}_3 : 532, \quad \mathbf{z}_4 : 2, \quad \mathbf{z}_5 : 1, \quad \mathbf{z}_{48} : 1, \quad \mathbf{z}_{49} : 1,$$

so 8 of the 50 candidates fire and 42 of the 49 tuples never do. Separately, of the 107 dual-feasibility polynomials, *all* 107 pass the Bernstein test at the root of the pipeline with no subdivision at all; the maximum subdivision depth over the whole feasibility half of the certificate is 0. The subdivision loop that [1] describes, and whose termination it can only assert, is needed for the cover alone.

Minimality of the four-tuple set. Re-running the cover search on subsets of the candidate list, at a budget of depth 40 and 30000 boxes, gives:

candidates	$x \in [0, 1/4]$	$x \in [1/4, 1/2]$	$x \in [1/2, 1]$
$\mathbf{z}_1, \mathbf{z}_2, \mathbf{z}_3, \mathbf{z}_{49}$	54	11	957
without $\mathbf{z}_1$	fails	fails	957
without $\mathbf{z}_2$	54	fails	fails
without $\mathbf{z}_3$	fails	fails	fails
without $\mathbf{z}_{49}$	54	11	fails
without $\mathbf{z}_{49}$, with (C3) restored	54	11	fails

Here "fails" means that the search exhausted its budget, and "$\mathbf{z}_2$" is always taken with its threshold $x \geq 1/4$. Two further searches: the cover also fails when all 48 tuples *other than* $\mathbf{z}_2$ are offered with no threshold restriction, and when $\mathbf{z}_2$ is restricted to $x \geq 1/2$. Thresholds $x \geq 1/8$ and $x \geq 1/4$ for $\mathbf{z}_2$ both close with the full candidate list (1014 and 1019 boxes); $x \geq 1/2$ does not.

One caveat is essential. A failure of the Bernstein test is not a proof that the disjunction is false: the test is sufficient for nonnegativity but not necessary, and the search is a search. What the table shows is that no accepted subdivision within this budget exists for the smaller candidate sets, which is the sense in which the four-tuple certificate is minimal. One instance of it is genuine and machine-checked, namely that the specific accepted 957-box subdivision of $x \in [1/2, 1]$ is rejected when only $\mathbf{z}_1, \mathbf{z}_3, \mathbf{z}_{49}$ are offered (Proposition 7.1(iii)).

Independent confirmation. The reduced certificate is expressible in the published verifier's own input schema, by setting the tangency parameter list to $\{60/100\}$ and the family-2 threshold to $1/4$. On that input the verifier prints its acceptance message with no edit to its code, in 3 min 35 s of wall time and 119 CPU-s, a factor 2.9 faster than on the published candidate. Its schema validator independently re-derives p from $D^\star$ and re-checks that T is consistent with the rule mixture, so the reduction cannot have substituted a different rule or a weaker bound.

7. Controls

If every check inside a certificate succeeds, nothing has been learned about the checks themselves. The following are machine-checked in the same way as the theorems above, and all of them are negative.

Proposition 7.1. (i) The constant is load-bearing. *Lower the target from T to the one that would give distortion 2.3 — that is, replace $\kappa = 6641/10000$ by $13/20$, decreasing $T(q)$ by $705q/42059$ — and recompute the four outer polynomials. At the explicit rational point $(q, x) = (3/8, 9/16)$ of the unit square, all four are strictly negative. So no candidate of the reduced set applies there, and Theorem 5.1 has no analogue at that constant.*

(ii) The slab hypothesis of Theorem 4.2 is not vacuous. *$H^{\mathrm{feas}}_{5,2}$ fails the root Bernstein test on the whole unit box, and is strictly negative at the explicit point $(x, y) = (1/16, 15/16)$.*

(iii) $\mathbf{z}_2$ is load-bearing in the cover. *The accepted 957-box subdivision of $x \in [1/2, 1]$ is rejected when offered only $\mathbf{z}_1, \mathbf{z}_3, \mathbf{z}_{49}$.*

(iv) The exact-division identities are not vacuous. *Perturbing one quotient by the constant 1 makes its identity fail.*

(v) The strict margin is sharp. $H^{\mathrm{feas}}_{1,49} \geq 1$ *is certified, and the same test one unit above the true minimum Bernstein coefficient is rejected.*

8. What is not claimed

The source's chain from Lemma 5 to $\mathrm{dist}(P^\star) \leq 11641/5000$ has four links that are not formalized here. None of them is a computation; all four are ordinary mathematics that a proof assistant can express but not cheaply.

(1) **The set-certificate proposition** (Proposition 2.1), due to [4] and restated by [5]: $\ell(P, I) \leq \kappa q(I)$ for every I implies $\mathrm{dist}(P) \leq 1 + 2\kappa$. This is the cheapest of the four and the only one that is purely finite-dimensional once $q(I)$ and $\ell(P, I)$ are defined on a finite profile.

(2) **The loss-envelope lemma**, which bounds $\ell(\mathrm{RSL}(D), I)$ by the value of an optimization over a scalar and two Borel measures — the measure-theoretic heart of [1], and the step that abstracts a profile away.

(3) **Weak duality for generalized moment problems** [12] together with the *sign chain* of [1, App. B.3]: from "$H^{\mathrm{feas}}_{1,j} > 0$ and $H^{\mathrm{feas}}_{2,j}, \ldots, H^{\mathrm{feas}}_{6,j} \geq 0$" to "$\mathbf{z}_j$ is dual-feasible". The source argues this in prose, from convexity of $G_x(\cdot, y)$ in b_0 — which reduces the first pointwise dual constraint to its two endpoint values — and from convexity of $v + \lambda f(xb_1) + \mu F_x(b_1) - b_1$ in b_1. Theorems 4.1–4.4 prove the polynomial *signs*; they do not prove dual feasibility.

(4) **Existence of** $\mathrm{RSL}(D^\star)$, a convex minimax argument, together with the endpoint cases $x = 0$ and $x = 1$ — including the primal-infeasibility argument at $x = 1$ for $q < 1 - p$, which is what confines the second verified disjunction to $q \in [1 - p, 1]$.

Both convexity facts in item 3 are statements about explicit polynomials with nonnegative coefficients in the right variables, and look reachable with the same instrument; that would upgrade Theorems 4.1–4.4 from sign theorems to feasibility theorems. We have not done it.

Remark 8.1 (the reach of the source's tuples). One more measurement, recorded so that it is not repriced. The 49 tuples do not depend on T; only the outer polynomials do. Write $M(q) = \sup_{x \in (0,1]} \min_j U(q, x; \mathbf{z}_j)$. The best target of the source's shape $T(q) = aq - bq^2$ then has $a(b) = \max_q \big(M(q)/q + bq\big)$, and the constant it certifies is $\kappa = a/(1 + b)$. A floating-point computation of M on a 300-point grid in q with four rounds of local refinement in x, followed by the one-parameter minimization of $a(b)/(1 + b)$, gives $\kappa \approx 0.66408958$ at $b \approx 0.1888$, i.e. a distortion of about 2.3281792; at the source's own $b = 7941/42059$ it gives $\kappa \approx 0.66408972$. So the published $\kappa = 0.6641$ carries about 1.04×10^{-5} of slack over what its own tuple family permits, and the available improvement is about 2.1×10^{-5} in the distortion — in the fifth decimal. This is consistent with the source's own discussion, which estimates its numerical search at "approximately 2.32818". The estimate is floating point and is not a theorem. Certifying past it was priced and not attempted: the box count of the cover scales roughly inversely with the slack, so a certified $\kappa = 0.66409$ would need on the order of 25 times the 1022 boxes, about 2.7 CPU-hours and 8 GB, for a fifth-decimal gain.

9. Verification

Theorems 4.1–4.4, Theorem 5.1, Proposition 5.2 and Proposition 7.1 have been formally verified in Lean 4 (toolchain `v4.32.0`) against *Mathlib* [13, 14]; the development is five modules and about 1770 lines, contains no `sorry` and declares no axiom by hand, and the whole chain checks cold in about 12 minutes of wall time and 12 CPU-minutes at a peak of 7.8 GB resident, of which the cover alone is 6 min 31 s and 577 CPU-s. The reasoning layer is free of compiled evaluation: the soundness of the Bernstein test, the new slab lemma — nonnegativity on a dyadic sub-interval of one axis, built from the two existing half-box substitutions and one affine reparametrization — and the new disjunctive tree walk, whose conclusion is that at every point of the box some member of the candidate list the checker was *started* with is nonnegative there, all depend on no axioms beyond `propext`, `Classical.choice` and `Quot.sound`. What is delegated to compiled evaluation, through Lean's `native_decide`, is exactly 56 Boolean tests of the form "this computed integer coefficient list

has no negative entry" or "this computed term list is empty": 40 for the feasibility half (a root test and an exact-division test per polynomial, with two slab tests instead of one for each of dual family 2's six), 3 subdivision-tree tests and 4 exact-division tests for the cover, and 9 for the controls, five of which are point witnesses that must return *true* and four of which are tests that must return *false*. Each theorem depends on precisely the Booleans its own proof uses, and those 56 Booleans are the only source of extra axioms anywhere in this paper. No coefficient list is supplied as an input to a sign claim: the certificate polynomials are built inside the proof from f, f' and Ψ by the identities of Lemma 3.1, so the Bernstein coefficients are the output of the pipeline; the exact quotients, which that algebra cannot build, are carried as coefficient data and are pinned by the division identities stated in Theorems 4.1–4.4 and Proposition 5.2, so a wrong coefficient makes an identity fail rather than a theorem weaker. Finally, every numerical value quoted in §6 and Remark 8.1 was produced by an independent implementation, and the reduction was confirmed by the source's own exact-rational verifier before any of it was formalized. The source of the development is currently held in a private repository: repository reference to be added at submission.

References

[1] N. Shah, *Improving randomized metric distortion to* 2.3282, arXiv:2608.29308v1 [cs.GT], 29 August 2026. All lemma, table and appendix numbers cited here are those of v1, read from the LaTeX source of the arXiv submission and resolved by compiling it; checked live on 3 September 2026, v1 is the only version. Theorem 1 is the bound 11641/5000; Lemma 5 is the inequality $g_{D^\star}(q) \le q(a - bq)$ whose computational half is the subject of this note; Appendix B.2 gives the dual tuples (Table 3) and their eligibility ranges (Table 4), Appendix B.3 the H^{feas} polynomials and the sign chain, and Appendix B.4 the outer polynomials, the verification method and the proof of Lemma 5.

[2] The verification code accompanying [1], `github.com/nisarg89/metric-distortion`: a standalone exact-rational verifier in the Python standard library together with the rational candidate data. Two commits, both dated 29 August 2026, unchanged as of 3 September 2026. Its verdict was reproduced on the published candidate before any claim in this note was made, and it accepts the reduced candidate of §6.

[3] J. Garloff, *The Bernstein algorithm*, Interval Computations **1993**, no. 2, 154–168 (Zbl 0829.65017). Not consulted directly; used through the full statement of the test given in [1, App. B.4].

[4] M. Charikar, P. Ramakrishnan, K. Wang and H. Wu, *Breaking the metric voting distortion barrier*, Journal of the ACM **71** (2024), no. 6, art. 42, 33 pp., doi:10.1145/3689625. Conference version: SODA 2024, 1621–1640, doi:10.1137/1.9781611977912.65; preprint arXiv:2306.17838.

[5] F. Frank, *An improved bound for the randomized metric distortion problem*, arXiv:2608.17863v1 [cs.GT], 18 August 2026.

[6] Q. Ye, *Half veto, half maximal lottery, five-halves distortion*, arXiv:2608.21202v1 [cs.GT], 21 August 2026.

[7] M. Charikar, P. Ramakrishnan, Z. Tan and K. Wang, *Metric distortion for tournament voting and beyond*, in: Proceedings of the 26th ACM Conference on Economics and Computation (EC 2025), 790–818, doi:10.1145/3736252.3742629.

[8] V. Gkatzelis, D. Halpern and N. Shah, *Resolving the optimal metric distortion conjecture*, in: Proceedings of the 61st IEEE Symposium on Foundations of Computer Science (FOCS 2020), 1427–1438, doi:10.1109/FOCS46700.2020.00134.

[9] E. Anshelevich, O. Bhardwaj, E. Elkind, J. Postl and P. Skowron, *Approximating optimal social choice under metric preferences*, Artificial Intelligence **264** (2018), 27–51, doi:10.1016/j.artint.2018.07.006.

[10] M. Charikar and P. Ramakrishnan, *Metric distortion bounds for randomized social choice*, in: Proceedings of the 33rd ACM-SIAM Symposium on Discrete Algorithms (SODA 2022), 2986–3004, doi:10.1137/1.9781611977073.116.

[11] F. E. Kizilkaya and D. Kempe, *Generalized veto core and a practical voting rule with optimal metric distortion*, in: Proceedings of the 24th ACM Conference on Economics and Computation (EC 2023), 913–936, doi:10.1145/3580507.3597798. The Simultaneous Plurality Veto process underlying Integrated Veto.

[12] A. Shapiro, *On duality theory of conic linear problems*, in: Semi-Infinite Programming: Recent Advances, Nonconvex Optimization and Its Applications, Springer, 2001, 135–165, doi:10.1007/978-1-4757-3403-4_7. Corollary 3.1 is the generalized-moment duality invoked by [1, Lemma 4].

[13] L. de Moura and S. Ullrich, *The Lean 4 theorem prover and programming language*, in: Automated Deduction — CADE 28, Lecture Notes in Computer Science **12699**, Springer, 2021, 625–635, doi:10.1007/978-3-030-79876-5_37.

[14] The mathlib Community, *The Lean mathematical library*, in: Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2020), ACM, 2020, 367–381, doi:10.1145/3372885.3373824.