THE SIMPLIFIED DUAL OF A CONDORCET-DIMENSION PROGRAM HAS OPTIMUM EXACTLY $1/(k+1)$: A CONJECTURE TRUE WITH ROOM TO SPARE, AND THE WEAK-DUALITY STEP THAT DOES NOT CLOSE

KOREA SUPERINTELLIGENCE LABS

Abstract. Let C be a set of m candidates, $\mathcal{S}$ the $m!$ strict rankings of C, and $\mathcal{M}_k$ the committees of size k. Zilberstein, Berker, Li and Martins [1] attack the question whether every election admits a Condorcet winning set of size 4 with a mixed-integer program, and then propose an analytic route: relax that program, simplify its dual to

$$\min u \quad \text{s.t.} \quad \sum \gamma_{W,c} = 1, \qquad u \geq \sum_{c \succ_s W} \gamma_{W,c} \quad (s \in \mathcal{S}), \qquad \gamma \geq 0,$$

and conjecture that its optimum u_k^* satisfies $u_k^* \leq 2/k$, whence — "by weak duality" — the bound $\alpha \leq 2/k$ and Condorcet winning sets of size 4. We evaluate the program. Read as that paper's own probabilistic restatement of the conjecture has it, with the pairs ranging over $\{(W,c) : W \in \mathcal{M}_k,\ c \in C \setminus W\}$, its optimum is exactly $1/(k+1)$ for every $m > k \geq 0$, independent of m. So the conjecture is true — but by a factor $2(k+1)/k > 2$, so it cannot be tight; read literally, with the printed index set $c \in C$, the optimum is 0 and the conjecture is vacuous. The inference it is meant to support does not hold: explicit finite elections at $k = 1, 2, 3, 4$ — two of them printed in [1], the other three attaining values its own table prints — have $\alpha^*(m,k)$ strictly above $1/(k+1)$. We locate the gap in the derivation — the step eliminating the big-M constant Q needs a quantity that is always at most $1/(m-k)$ to equal 1 — and we close the route at its root: the relaxation being dualised has optimum exactly $1/(k+1) + Q(m-k-1)/(m-k)$, which at the paper's own $Q = 1$ is at least 1 for all $m \geq 2k+1$; so every bound weak duality against it can certify is at least 1, and for $k \geq 3$ — in particular for the $k = 4$ the conclusion turns on — no dual certificate for it can prove $\alpha \leq 2/k$. Nothing here bounds the Condorcet dimension: whether every election has a Condorcet winning set of size 4 remains open. Every statement below is machine-checked in Lean 4, apart from those that say, at the point where they stand, that they are external computations.

1. Introduction

The objects. Fix a finite set C of m *candidates*. A *ranking* is a strict linear order on C; we record it as a bijection $s \colon C \to \{0, 1, \ldots, m-1\}$ assigning each candidate its position, 0 being best, and write $c \succ_s c'$ for $s(c) < s(c')$. Let $\mathcal{S}$ be the set of all $m!$ rankings. An *election*, or *profile*, is a probability distribution $x \in \Delta(\mathcal{S})$ over rankings; a finite list of voters with a ranking each is the special case of rational x with denominator the number of voters.

A *committee* is a subset $W \subseteq C$, and $\mathcal{M}_k$ denotes the set of committees of size k. For a ranking s, a candidate c and a committee W we write

$$c \succ_s W \quad :\Longleftrightarrow \quad s(c) < s(w) \ \text{ for every } w \in W,$$

read "c beats W under s": the voter would rather have c than any member of W. The support of the challenge (W,c) under a profile x is $\mathrm{spt}_x(W,c) = \sum_{s\,:\,c \succ_s W} x_s$.

Following Elkind, Lang and Saffidine [2], a committee W is α-*undominated* in x when $\mathrm{spt}_x(W,c) < \alpha$ for every $c \notin W$, and a *Condorcet winning set* is a $\frac{1}{2}$-undominated committee: no candidate outside it is preferred to all of it by a majority. The *Condorcet dimension* of an election is the size of its smallest Condorcet winning set. The governing quantity is the min–max

$$\alpha^*(m,k) \ = \max_{x \in \Delta(\mathcal{S})} \ \min_{W \in \mathcal{M}_k} \ \max_{c \notin W} \ \mathrm{spt}_x(W,c), \tag{1}$$

which [1] prints as its equation (1): if $\alpha^*(m,k) < \frac{1}{2}$ for every m, then every election has a Condorcet winning set of size k.

The state of the question. Every election has a Condorcet winning set of constant size. Charikar, Lassota, Ramakrishnan, Vetta and Wang [3] proved that six candidates suffice, and Song, Nguyen and Lin [4] brought the bound to 5; on the other side, [2] exhibits elections of Condorcet dimension 3. The gap $3 \le \cdot \le 5$ has not moved: a paper posted in August 2026 states the position, in a footnote of its introduction, as

> "It is known that every election has Condorcet winning sets of constant size, even as small as 5. The possibility of an improvement to 4 or 3 remains a fascinating open problem." [5]

(the citation markers carried by that sentence are elided here). The present paper does not narrow that gap and does not attempt to; see Section 11.

The program under examination. Zilberstein, Berker, Li and Martins [1] search for elections of large Condorcet dimension with a mixed-integer program, do not beat the known lower bound of 3, and then offer what they call a "concrete analytical path" to an upper bound. They encode (1) as a MILP with a big-M constant Q and binary challenger selectors $y_{W,c}$, take its LP relaxation, dualise, delete the terms carrying Q, and arrive at a program they call the *simplified dual*, whose optimum they call u_k^*. Their conjecture, and the use they propose for it, are quoted in full in Section 2; in outline:

> "We conjecture that for any number of candidates m and committee size k, this value is bounded by: $u_k^* \le \frac{2}{k}$" …"By weak duality, proving this conjecture would immediately imply that $\alpha \le \frac{2}{k}$, meaning that Condorcet winning sets of size 4 always exist." [1], Conjecture 5.1

This is an inviting target: the simplified dual is a small, explicit, highly symmetric linear program, and nothing about it is conjectural once one computes it. That is what we do.

What this paper settles.

(1) **The value** (Theorem 4.1). Read as the paper's own restatement of its conjecture — the "Dual Bound Probabilistic Interpretation", whose pairs range over $\{(W,c) : W \in \mathcal{M}_k,\ c \in C \setminus W\}$ — the simplified dual has optimum

$$u_{m,k}^* = \frac{1}{k+1} \qquad \text{for every } m > k \ge 0,$$

exactly, and in particular independently of the number of candidates m. Two counting lemmas do it: each ranking covers exactly $\binom{m}{k+1}$ of the $(m-k)\binom{m}{k}$ pairs, and each pair is covered by exactly $m!/(k+1)$ of the $m!$ rankings. The first gives the uniform γ as an optimal solution; the second gives the matching lower bound by averaging, and says that no distribution of challenges beats the uniform one.

(2) **The conjecture, settled** (Corollary 4.2). Since $1/(k+1) \le 2/k$ for all $k \ge 1$, the conjecture of [1] is *true*. It is true by the factor $2(k+1)/k > 2$, so it is not tight, and the quantity it bounds is nowhere near the bound proposed for it.

(3) **The printed program is vacuous** (Theorem 5.1). The constraint as printed sums over all $c \in C$, not over $c \in C \setminus W$. A pair with $c \in W$ is never covered, since $c \succ_s W$ would require $s(c) < s(c)$; parking all the mass on such a pair gives optimum 0. The paper's two statements of its own conjecture therefore have different optima, 0 and $1/(k+1)$, and we prove both so that the ambiguity is resolved rather than assumed away.

(4) **Weak duality against this program fails** (Theorem 6.1). For $k = 1, 2, 3, 4$ there are explicit finite elections in which *no* committee of size k is α-undominated for an α strictly larger than $u_{m,k}^* = 1/(k+1)$. Every witness is tied to [1] itself: its opening Condorcet-paradox table (Table 1), its Figure 1, and profiles attaining three of the values its own MILP table prints. The case $k = 1$ is the sharpest: $u_{m,1}^* = \frac{1}{2}$ for every m, so weak duality against the simplified dual would assert that every election has a candidate whom fewer than half the voters wish to replace — that Condorcet's paradox does not occur.

(5) **Where the derivation loses the bound** (Theorem 7.1). The step that eliminates Q requires $\sum_{W,c}\gamma_{W,c} - \sum_W \min_c \gamma_{W,c}$ to vanish. It cannot: a minimum is at most an average, so $\sum_W \min_c \gamma_{W,c} \le 1/(m-k)$ for every feasible γ, and the coefficient of Q is at least $1 - 1/(m-k) \ge \frac{1}{2}$ as soon as $m \ge k+2$. The deleted term is not a rounding error; at $(m,k,Q) = (4,2,1)$ it is exactly the difference $5/6 - 1/3$ between the relaxation's optimum and the simplified dual's.

(6) **The route is closed at its root** (Theorem 8.1 and Corollary 8.2). The LP relaxation being dualised has optimum exactly

$$\mathrm{LP}^*(m,k,Q) = \frac{1}{k+1} + Q\,\frac{m-k-1}{m-k},$$

for every $m > k \ge 0$ and every $Q \ge 0$. At the paper's own $Q = 1$ this is at least 1 whenever $m \ge 2k+1$, and $2/k \le 1$ for every $k \ge 2$; so the relaxation is already weaker than the trivial bound $\alpha \le 1$, and for $k \ge 3$ — the case $k = 4$ included — *no* dual certificate for it can establish $\alpha \le 2/k$. This is independent of item (5): it survives keeping the Q terms and dualising correctly. It also turns into a theorem the numerical observation [1] records in the caption of its Table 6, that "as m increases, the relaxation appears to approach 1", and sharpens it: the value does not approach 1, it passes 1, at $m = 2k+1$.

(7) **Which program the printed table is** (Proposition 9.1). Table 6 of [1], captioned "Optimal values for the LP relaxation of `MILP`", is matched in all 21 of its cells by the relaxation *with* that paper's "Fixed W^*" symmetry breaking, and disagrees with the plain relaxation of item (6) at every one of its fifteen off-diagonal cells: at $(m,k) = (4,2)$ the plain relaxation is $5/6$ where the table prints 0.750.

Items (1), (2) and (6) are offered as new; items (3), (4), (5) and (7) are corrections to the derivation and the presentation of [1], and Section 12 says which is which and what search stands behind each.

On the corrections, and on sources. [1] is a workshop paper of April 2026 whose authors are active, and the errata below are offered in that spirit. They are also, we think, good news for it: the conjecture it states turns out to be true and to have a four-line proof, and the obstruction we exhibit is at the level of a modelling choice — the big-M relaxation — not of the idea of certifying α by duality.

All quotations are from arXiv:2604.19851v1, read from the LaTeX e-print on 2026-09-02 and rechecked on 2026-09-03; version 1 is the only version. Every equation, table, figure and conjecture of [1] named below is named by the number it carries in the compiled v1 document, obtained by compiling that e-print: equation (1) is α^*; equations (2)–(5) are the MILP; equation (11) is the dual constraint $\beta_W + Q\gamma_{W,c} \ge 0$; equations (14)–(15) are the simplified dual; Table 1 is the opening Condorcet-paradox election, Table 3 the MILP values and Table 6 the LP-relaxation values; Figure 1 is the three-voter example election; and the conjecture is Conjecture 5.1, with Conjecture 5.2 its probabilistic restatement. (The same conjecture is previewed as Conjecture 1.1 in the introduction of [1], with "defined in Section 5" in place of the constraint numbers.) Since these numbers could move if a version 2 appeared, each is pinned to v1.

2. The three programs

Throughout, $m > k \ge 0$ unless stated otherwise, and all arithmetic is over $\mathbb{Q}$.

Definition 2.1 (pairs and coverage)**.** The set of *committee–challenger pairs* is

$$\mathcal{P}_{m,k} = \{(W,c)\ :\ W \in \mathcal{M}_k,\ c \in C \setminus W\}, \qquad |\mathcal{P}_{m,k}| = \binom{m}{k}(m-k),$$

and for a ranking $s \in \mathcal{S}$ the pairs it *covers* are $\mathrm{Cov}_{m,k}(s) = \{(W,c) \in \mathcal{P}_{m,k} : c \succ_s W\}$.

The MILP and its relaxation. The encoding of (1) in [1] carries a variable $x_s \geq 0$ for each ranking, a selector $y_{W,c}$ for each pair, a big-M constant Q, and the objective α:

$$\begin{aligned} \max\ & \alpha \\ \text{s.t.}\ & \textstyle\sum_{s \in \mathcal{S}} x_s = 1, \qquad x \geq 0, \\ & \textstyle\sum_{c \notin W} y_{W,c} = 1 \quad (W \in \mathcal{M}_k), \qquad y \geq 0, \\ & \operatorname{spt}_x(W,c) + Q\,(1 - y_{W,c}) \ \geq\ \alpha \quad ((W,c) \in \mathcal{P}_{m,k}). \end{aligned} \tag{2}$$

With y binary, the selector picks for each committee the challenger whose constraint binds; the LP relaxation (2) lets y be fractional. (The bound $y \leq 1$ is implied by $y \geq 0$ and the selector equation, so we do not carry it.) Write $\mathrm{LP}^*(m,k,Q)$ for the optimum of (2). The paper sets $Q = 1$.

The simplified dual. Dualising (2) produces multipliers $\gamma_{W,c}$ for the last family of constraints and β_W for the selector equations. [1] eliminates β and then argues that the terms carrying Q vanish; the program it is left with is printed there as

> "$\min u$ s.t. $\sum_{W \in \mathcal{M}_k, c \in C} \gamma_{W,c} = 1$; $u \geq \sum_{W \in \mathcal{M}_k,\, c \succ_s W} \gamma_{W,c}$, $\forall s \in \mathcal{S}$; $\gamma_{W,c} \geq 0$" [1], equations (14)–(15)

and its optimum is called u_k^*. The conjecture is

> "**Conjecture 5.1 (Dual Bound on α).** Let u_k^* denote the optimal objective value of the simplified dual linear program defined by constraints (14)–(15). We conjecture that for any number of candidates m and committee size k, this value is bounded by: $u_k^* \leq \frac{2}{k}$",

followed by

> "By weak duality, proving this conjecture would immediately imply that $\alpha \leq \frac{2}{k}$, meaning that Condorcet winning sets of size 4 always exist." [1]

The same paper restates the conjecture in words, and the restatement fixes an index set that the displayed program leaves ambiguous:

> "**Conjecture 5.2 (Dual Bound Probabilistic Interpretation).** There exists a probability distribution γ over the set of committee-challenger pairs $\{(W,c) \mid W \in \mathcal{M}_k,\ c \in C \setminus W\}$ such that for every possible ranking $s \in \mathcal{S}$, the total probability mass of pairs where the challenger defeats the committee is at most $2/k$." [1]

The displayed constraint ranges over $c \in C$; the restatement ranges over $c \in C \setminus W$. We treat both.

Definition 2.2 (the two readings)**.** Call $\gamma\colon \mathcal{P}_{m,k} \to \mathbb{Q}$ *feasible* when $\gamma \geq 0$ and $\sum_{p \in \mathcal{P}_{m,k}} \gamma_p = 1$, and put $\operatorname{val}(\gamma,s) = \sum_{p \in \operatorname{Cov}_{m,k}(s)} \gamma_p$. A number v is *the optimum of the simplified dual*, written $u_{m,k}^* = v$, when

1. some feasible γ has $\operatorname{val}(\gamma,s) \leq v$ for every $s \in \mathcal{S}$, and
2. every feasible γ has $\operatorname{val}(\gamma,s) \geq v$ for some $s \in \mathcal{S}$.

The *literal* reading replaces $\mathcal{P}_{m,k}$ throughout by $\mathcal{P}_{m,k}^{\mathrm{lit}} = \mathcal{M}_k \times C$, and its optimum is written $u_{m,k}^{*,\mathrm{lit}}$. Both notions pin at most one value (Proposition 10.1 (4)).

Conditions (1) and (2) together say exactly that the minimum over feasible γ of $\max_s \operatorname{val}(\gamma,s)$ is attained and equals v; we state them separately because that is the form in which each is used and proved.

3. Two counting lemmas

Both lemmas are exact counts, uniform in the argument, and everything else in the paper is built on them.

Lemma 3.1 (each ranking covers $\binom{m}{k+1}$ pairs)**.** *For every m, every k and every ranking $s \in \mathcal{S}$,*

$$|\operatorname{Cov}_{m,k}(s)| = \binom{m}{k+1},$$

independently of s.

Proof. The map $(W, c) \mapsto W \cup \{c\}$ sends $\mathrm{Cov}_{m,k}(s)$ to the $(k+1)$-subsets of C: it lands there because $c \notin W$. It is injective: if (W, c) and (W', c') are covered pairs with $W \cup \{c\} = W' \cup \{c'\}$ and $c \neq c'$, then $c' \in W$ and $c \in W'$, so $s(c) < s(c')$ and $s(c') < s(c)$; hence $c = c'$, and erasing c from the common set gives $W = W'$. It is surjective: a $(k+1)$-set T is nonempty, so it has a unique s-best element c, and $(T \setminus \{c\}, c)$ is a covered pair mapping to T. The $(k+1)$-subsets number $\binom{m}{k+1}$. □

Lemma 3.2 (each pair is covered by $m!/(k+1)$ rankings)**.** *For every committee W with $|W| = k$ and every $c \notin W$,*

$$\bigl|\{s \in \mathcal{S} \ : \ c \succ_s W\}\bigr| \cdot (k+1) = m!.$$

Proof. Put $T = W \cup \{c\}$, so $|T| = k+1$, and for $u \in T$ let A_u be the set of rankings in which u is the s-best element of T. Every ranking makes exactly one element of the nonempty finite set T best, so the A_u, $u \in T$, partition $\mathcal{S}$. For $u, v \in T$ the map $s \mapsto s \circ (u\,v)$, composition with the transposition exchanging u and v, carries A_u into A_v and is an involution, so all $k+1$ classes have the same size. Finally $\{s : c \succ_s W\} = A_c$, whence $(k+1)|A_c| = |\mathcal{S}| = m!$. □

4. The optimum of the simplified dual

Theorem 4.1 (the value)**.** *For every $m > k \geq 0$,*

$$u^*_{m,k} = \frac{1}{k+1}.$$

In particular the optimum does not depend on the number of candidates m.

Proof. Upper bound. Let γ be uniform on $\mathcal{P}_{m,k}$, that is $\gamma_p = 1/|\mathcal{P}_{m,k}| = 1/\bigl(\binom{m}{k}(m-k)\bigr)$, which is feasible since $k < m$ makes $\mathcal{P}_{m,k}$ nonempty. By Lemma 3.1 every ranking has

$$\mathrm{val}(\gamma, s) = \frac{\binom{m}{k+1}}{\binom{m}{k}(m-k)} = \frac{1}{k+1},$$

the last step being the identity $\binom{m}{k+1}(k+1) = \binom{m}{k}(m-k)$. So condition (1) of Definition 2.2 holds at $v = 1/(k+1)$, with equality for every s.

Lower bound. Let γ be any feasible point. Summing $\mathrm{val}(\gamma, \cdot)$ over all $m!$ rankings and exchanging the order of summation,

$$\sum_{s \in \mathcal{S}} \mathrm{val}(\gamma, s) = \sum_{p \in \mathcal{P}_{m,k}} \gamma_p \cdot \bigl|\{s : s \text{ covers } p\}\bigr| = \frac{m!}{k+1} \sum_{p \in \mathcal{P}_{m,k}} \gamma_p = \frac{m!}{k+1},$$

by Lemma 3.2 and feasibility. A finite family of numbers cannot all lie strictly below their mean, so some s has $\mathrm{val}(\gamma, s) \geq 1/(k+1)$, which is condition (2). □

Neither half is an exhaustive computation: both are uniform in m and k, and no election, no committee and no ranking is enumerated anywhere in the argument.

Corollary 4.2 (the conjecture of [1] is true, with a factor-2 margin)**.** *For every $m > k \geq 1$ the Dual Bound conjecture holds:*

$$u^*_{m,k} = \frac{1}{k+1} < \frac{2}{k}, \qquad \textit{and indeed} \qquad 2\,u^*_{m,k} \leq \frac{2}{k}.$$

Proof. Immediate from Theorem 4.1: $k < 2(k+1)$ gives the strict inequality, and $2k \leq 2(k+1)$ the second. □

The margin is the substance of Corollary 4.2. The conjectured bound exceeds the truth by the factor $2(k+1)/k$, which is > 2 for every $k \geq 1$ and tends to 2; a bound true by more than a factor of two is not the bound that was being conjectured, and Section 6 shows that the true value is not an upper bound for α^* either.

Remark 4.3 (the constant is classical; its optimality is not). The number $1/(k+1)$ is the impartial-culture probability that a uniformly random ranking puts a fixed candidate above all k members of a fixed committee, and it appears in this literature as the large-voter threshold $1 - 1/(k+1)$ for winning committees, credited there to [2], Corollary 2; see [6]. What Theorem 4.1 adds is that this is the *exact* optimum of the program of [1] — in particular the matching lower bound, that no distribution of challenges does better than the uniform one, which is the half that makes Corollary 4.2 and Section 6 possible.

5. The printed program, read literally

Theorem 5.1 (the literal reading is vacuous). *Let $1 \le k \le m$. Then $u^{*,\mathrm{lit}}_{m,k} = 0$: there is a distribution on $\mathcal{M}_k \times C$ no ranking covers at all. Consequently, for $m > k \ge 1$ the two readings of the conjecture of* [1] *have different optima,*

$$u^{*,\mathrm{lit}}_{m,k} = 0 \;\ne\; \frac{1}{k+1} = u^*_{m,k},$$

and the conjectured bound $u^*_k \le 2/k$ is vacuously true under the printed one.

Proof. Choose any $W \in \mathcal{M}_k$ and any $c \in W$; this is possible because $k \le m$ and $k \ge 1$. Put all the mass on the single pair (W, c). No ranking covers it: $c \succ_s W$ would require $s(c) < s(c)$. So every ranking has covered mass 0, which is condition (1) of Definition 2.2 at $v = 0$; condition (2) is trivial, since covered mass is a sum of nonnegative terms. The last sentence follows because $0 \ne 1/(k+1)$ and each reading pins at most one value. □

We stress that this is an editorial ambiguity, not a mathematical dispute: the paper's own probabilistic restatement of its conjecture uses $c \in C \setminus W$, and everything else in the present paper is proved under that reading. Theorem 5.1 is recorded so that a reader who takes the displayed constraint at face value knows which of the two programs is being discussed, and it is worth recording because the two are not close: one has optimum 0 and the other $1/(k+1)$.

6. Weak duality against the simplified dual

The inference proposed in [1] is that $u^*_{m,k}$ bounds $\alpha^*(m,k)$ from above. It does not, and every witness below comes from that paper's own data.

The five elections below are given by their rankings with integer voter counts, so that the verification is arithmetic over $\mathbb{N}$; we describe them, and refer to the formal development for the position vectors. $E_{6,1}$ is the six-voter, six-candidate cyclic election of [1], Table 1, captioned there "An election showing that for any candidate chosen as the winner, there are 5/6 voters who prefer some other candidate"; $E_{3,1}$ is the three-voter cycle of its Figure 1; and $E_{4,2}$, $E_{5,3}$, $E_{6,4}$ are profiles on 5, 10 and 13 voters whose values 2/5, 3/10 and 3/13 are the numbers its Table 3 — "Max α values from `MILP` with Fixed W^*, no subsampling, and 720s timeout" — prints as 0.400, 0.300 and 0.231 at $(m,k) = (4,2)$, $(5,3)$, $(6,4)$. Only the first two elections are printed in [1]; the last three were extracted here from a reproduction of those three cells (Section 12) and then re-verified from scratch.

Theorem 6.1 (weak duality fails at $k = 1,2,3,4$). *In each of the following, no committee of size k is α-undominated for the stated α, so $\alpha^*(m,k) \ge \alpha > u^*_{m,k}$:*

election	(m,k)	*voters*	α	$u^*_{m,k}$
$E_{6,1}$	$(6,1)$	6	$5/6$	$1/2$
$E_{3,1}$	$(3,1)$	3	$2/3$	$1/2$
$E_{4,2}$	$(4,2)$	5	$2/5$	$1/3$
$E_{5,3}$	$(5,3)$	10	$3/10$	$1/4$
$E_{6,4}$	$(6,4)$	13	$3/13$	$1/5$

In particular the optimum of the simplified dual is not an upper bound for α^*, at any of $k = 1,2,3,4$.

Proof. The values of $u^*_{m,k}$ are Theorem 4.1 at $(m, k) = (6, 1), (3, 1), (4, 2), (5, 3), (6, 4)$, and the five numerical inequalities $1/2 < 5/6$, $1/2 < 2/3$, $1/3 < 2/5$, $1/4 < 3/10$, $1/5 < 3/13$ are arithmetic. The claim about each election is an exhaustive finite check, and it is the only kind of exhaustive check in this paper: for the stated election E and every one of the $\binom{m}{k}$ committees W of size k — at most $\binom{6}{4} = 15$ of them — one exhibits a challenger $c \notin W$ with

$$a \cdot (\text{number of voters}) \ \le\ b \cdot (\text{number of voters ranking } c \text{ above all of } W),$$

where $\alpha = a/b$. Both sides are natural numbers, and the voter counts are sums over the at most 13 listed rankings, so the whole check is decided by evaluation in $\mathbb{N}$; no rational arithmetic and no floating point enters. Since this holds for every W, the profile x of E has $\min_W \max_{c\notin W} \operatorname{spt}_x(W, c) \ge a/b$, and $\alpha^*(m, k) \ge a/b$ by (1). □

The case $k = 1$ deserves a sentence. Theorem 4.1 gives $u^*_{m,1} = \frac{1}{2}$ for *every* m; so weak duality against the simplified dual would assert that every election has a candidate whom fewer than half of the voters wish to replace. That is the assertion that Condorcet's paradox does not occur, and the counterexample is the very table with which [1] opens. The case $k = 4$ is the one its conclusion turns on — "Condorcet winning sets of size 4 always exist" — and there the dual optimum is $u^*_{m,4} = 1/5$ against the 0.231 that paper prints itself.

Remark 6.2 (the failure is not confined to $k \le 4$; outside the formal development). The following construction is *not* part of the machine-checked development, and is recorded as an ordinary argument. For $k \ge 1$ let $m = k(k+2)$ and let E_m be the cyclic election on $C = \mathbb{Z}_m$ with m voters, voter j ranking candidate c in position $(c - j) \bmod m$. Fix W with $|W| = k$ and $c \notin W$, put $d_w = (w - c) \bmod m$ and $D = \max_{w\in W} d_w$; writing $t = (c - j) \bmod m$ for the position of c under voter j, the candidate w sits at position $(t + d_w) \bmod m$, which exceeds t exactly when $t < m - d_w$. Hence $c \succ_j W$ for exactly $m - D$ voters. Next, for $c \in \mathbb{Z}_m$ let $V_c = \{c, c-1, \dots, c-k\}$, of size $k+1$; each element of $\mathbb{Z}_m$ lies in exactly $k+1$ of the V_c, so $\sum_c |V_c \cap W| = k(k+1) < m$ and some V_c misses W. For that c we get $D \le m - k - 1$, so at least $k + 1$ voters put c above all of W, and therefore

$$\alpha^*\big(k(k+2), k\big) \ \ge\ \frac{k+1}{k(k+2)} \ >\ \frac{1}{k+1} \ =\ u^*_{k(k+2),k},$$

the strict inequality because $(k + 1)^2 = k^2 + 2k + 1 > k^2 + 2k$. So weak duality against the simplified dual fails at every committee size, not only the four of Theorem 6.1. The value $(k + 1)/(k(k + 2))$ was checked to be attained, by exhaustive enumeration of the committees, for $k = 1, 2, 3, 4$, where it is $2/3$, $3/8$, $4/15$, $5/24$ against $1/(k + 1) = 1/2$, $1/3$, $1/4$, $1/5$; at $k = 1$ it is the classical Condorcet paradox. This argument is outside the formal development because it needs modular arithmetic that the rest of the development avoids; it is stated here for completeness and nothing else in the paper depends on it.

7. Where the elimination of Q loses the bound

The simplified dual is obtained from the dual of (2) by discarding the terms carrying Q. The step is

> "From constraint (11), and the fact that we are minimizing, the optimal β_W will be tightly bound as: $\beta_W = -Q \cdot \min_{c\in C} \gamma_{W,c}$. Substituting this into the objective function, the terms involving Q are: $Q\big(\sum_{W,c} \gamma_{W,c} - \sum_W \min_c \gamma_{W,c}\big)$. For the objective value to be independent of Q, this term must vanish. Since $\gamma_{W,c} \ge 0$, the sum is always greater than or equal to the minimum. Equality implies that for each committee W, $\gamma_{W,c}$ must be concentrated on a single candidate c (or set of challengers with equal weight), effectively mimicking the integral behavior of the primal $y_{W,c}$ variables." [1]

The identification of the tight β_W is correct. What cannot hold is the vanishing. Note first that the remedy the passage offers is self-defeating: concentrating $\gamma_{W,\cdot}$ on a single challenger makes $\min_c \gamma_{W,c} = 0$, so the second sum becomes 0 rather than matching the first, which is 1. The following says that nothing else works either.

Theorem 7.1 (the coefficient of Q is bounded away from 0)**.** *Let $m > k \geq 0$, let γ be feasible for the simplified dual, and let $\mu\colon \mathcal{M}_k \to \mathbb{Q}$ satisfy $\mu(W) \leq \gamma_{W,c}$ for every $W \in \mathcal{M}_k$ and every $c \notin W$. Then*

$$(m-k) \sum_{W \in \mathcal{M}_k} \mu(W) \ \leq\ 1.$$

Consequently, if $m \geq k+2$ then

$$\sum_{(W,c) \in \mathcal{P}_{m,k}} \gamma_{W,c} \ -\ \sum_{W \in \mathcal{M}_k} \mu(W) \ \geq\ \frac{1}{2},$$

so the coefficient of Q never vanishes.

Proof. Fix $W \in \mathcal{M}_k$. The set $C \setminus W$ has $m-k$ elements and $\mu(W)$ is a lower bound for $\gamma_{W,c}$ on all of it, so $(m-k)\mu(W) \leq \sum_{c \notin W} \gamma_{W,c}$. Summing over $W \in \mathcal{M}_k$ and using feasibility,

$$(m-k) \sum_{W} \mu(W) \ \leq\ \sum_{W} \sum_{c \notin W} \gamma_{W,c} = \sum_{p \in \mathcal{P}_{m,k}} \gamma_p = 1.$$

For the second claim, $m \geq k+2$ gives $m-k \geq 2$, so $\sum_W \mu(W) \leq 1/(m-k) \leq \frac{1}{2}$ when the sum is positive, and the bound is trivial when it is not; subtract from $\sum_p \gamma_p = 1$. □

Taking $\mu(W) = \min_{c \notin W} \gamma_{W,c}$ recovers the quantity of the quoted passage. We state the theorem for an arbitrary pointwise lower bound μ so that it is insensitive to how the minimum is taken: the passage minimises over all $c \in C$, and that minimum is in particular a lower bound for the challenger weights, so the conclusion applies under either reading of the index set.

The size of the deleted term is not marginal. At $(m,k) = (4,2)$ with $Q = 1$, the relaxation's optimum is 5/6 (Theorem 8.1) and the simplified dual's is 1/3 (Theorem 4.1); the difference 1/2 is exactly $Q\bigl(1 - 1/(m-k)\bigr)$, that is, exactly the term that was deleted.

8. The relaxation itself, and why the route cannot reach its destination

The obstruction of the previous section is a defect in one step. The following is stronger and independent of it: even with the Q terms kept and the dual taken correctly, weak duality against (2) cannot deliver the conclusion, because (2) is already too weak.

Theorem 8.1 (the relaxation in closed form)**.** *For every $m > k \geq 0$ and every $Q \geq 0$, the optimum of the LP relaxation (2) is*

$$\mathrm{LP}^*(m,k,Q) \ =\ \frac{1}{k+1} + Q\,\frac{m-k-1}{m-k}.$$

Both halves hold: every feasible point has α at most that value, and it is attained.

Proof. Upper bound. Let (x,y,α) be feasible. Sum the constraint $\alpha \leq \mathrm{spt}_x(W,c) + Q(1-y_{W,c})$ over all $\binom{m}{k}(m-k)$ pairs. On the right, the selector equations give $\sum_{p \in \mathcal{P}_{m,k}} y_p = \binom{m}{k}$, and Lemma 3.1 with $\sum_s x_s = 1$ gives

$$\sum_{p \in \mathcal{P}_{m,k}} \mathrm{spt}_x(p) = \sum_{s \in \mathcal{S}} x_s \,|\operatorname{Cov}_{m,k}(s)| = \binom{m}{k+1},$$

so that

$$\binom{m}{k}(m-k)\,\alpha \ \leq\ \binom{m}{k+1} + Q\Bigl(\binom{m}{k}(m-k) - \binom{m}{k}\Bigr).$$

Dividing by $\binom{m}{k}(m-k) > 0$ and using $\binom{m}{k+1}(k+1) = \binom{m}{k}(m-k)$ gives the bound.

Attainment. Take x uniform on all $m!$ rankings and $y \equiv 1/(m-k)$. Both normalisations hold, and by Lemma 3.2 every pair has support exactly $1/(k+1)$, so every constraint of the last family holds with equality at $\alpha = 1/(k+1) + Q\bigl(1 - 1/(m-k)\bigr)$, which is the stated value. □

Again neither half enumerates anything: the argument is one summation and one explicit point, uniform in m, k and Q. (The machine-checked form of Theorem 8.1 carries no sign condition on Q at all; we state it for $Q \geq 0$, which is the case the application needs.)

Corollary 8.2 (the route is closed at its root)**.** *Let $k \geq 2$ and $m \geq 2k+1$. Then, at the value $Q = 1$ used in* [1]*, the relaxation* (2) *has a feasible point with*

$$\alpha \;\geq\; 1 \;\geq\; \frac{2}{k}.$$

So the relaxation's own optimum is at least the trivial bound $\alpha \leq 1$; and for $k \geq 3$, where $2/k < 1$, no dual certificate for (2) *can establish $\alpha \leq 2/k$.*

Proof. By Theorem 8.1 the point of its attainment half has $\alpha = 1/(k+1) + (m-k-1)/(m-k)$, and

$$\frac{1}{k+1} + \frac{m-k-1}{m-k} \;\geq\; 1 \iff (k+1)(m-k-1) \;\geq\; k(m-k) \iff m \;\geq\; 2k+1.$$

Finally $2/k \leq 1$ for $k \geq 2$. □

A word on how much Corollary 8.2 says. Any bound on α that weak duality against (2) can certify is at least the optimum of (2), hence at least 1 once $m \geq 2k+1$; that is where the conclusion comes from. The case $k = 2$ the corollary leaves aside is not a gap in the argument: there the conjectured bound is $2/k = 1$, which is the trivial bound $\alpha \leq 1$ and certifies nothing about Condorcet winning sets, so the route yields nothing at $k = 2$ either. The case the conclusion of [1] turns on is $k = 4$.

Corollary 8.2 makes precise, and sharpens, an observation [1] records numerically in the caption of its Table 6: "as m increases, the relaxation appears to approach 1". The closed form shows the value does not approach 1 but passes it, with equality exactly at $m = 2k+1$ — there $1/(k+1)+k/(k+1) = 1$ — and continues to grow. [1] reads its observation as a reason to study the dual instead; the reading available after Theorem 8.1 is that the relaxation has already lost the information the certificate would have to carry. On the diagonal $m = k+1$, by contrast, the Q term $(m-k-1)/(m-k)$ vanishes identically, the relaxation is exact at $1/m$, and the simplified dual *is* the true dual; the failure documented in this paper is an off-diagonal phenomenon. See Proposition 10.1 (5) and the paragraph that follows it.

9. Which program the printed table is

[1] prints a table, its Table 6, captioned "Optimal values for the LP relaxation of `MILP`. Note that as m increases, the relaxation appears to approach 1.", with 21 cells for $3 \leq m \leq 8$ and $2 \leq k \leq 7$. Its values are smaller than Theorem 8.1 off the diagonal. The caption names the LP relaxation of `MILP` and nothing else, but the same paper describes, under "Symmetry breaking", a modified program in which a committee $W^* = \{c_1, \ldots, c_k\}$ is fixed as the minimiser and c_{k+1} as its challenger: W^* is dropped from the selector and α constraint families, and the equality $\alpha = \sum_{s\,:\,c_{k+1} \succ_s W^*} x_s$ is added as an anchor. That paper's own Table 2 records Fixed W^* as the optimisation it adopts for `MILP`, and the caption of its Table 3 names it explicitly. The identification below is an inference from numerical agreement, not from a statement in [1]: the LP relaxation of the symmetry-broken program matches all 21 printed cells — nineteen of them exactly on rounding, the other two to within one unit in the last printed digit — while the plain relaxation (2) of Theorem 8.1 matches only the six on the diagonal and misses all fifteen others.

Proposition 9.1 (the table is the symmetry-broken relaxation)**.** *At $(m,k) = (4,2)$ and $Q = 1$ the plain relaxation* (2) *has optimum $5/6$, while the printed cell is 0.750, and $3/4 < 5/6$. Nineteen of the 21 printed cells are the optimum of the relaxation with Fixed W^* rounded to three decimals; at $(7,4)$ the printed 0.777 is that optimum, $7/9$, truncated instead of rounded, and at $(7,2)$ the optimum is 1 where the table prints 0.999. The plain relaxation agrees with the printed table only on the diagonal $k = m-1$, where both equal $1/m$, and differs from it at all fifteen other cells.*

Proof and computational status. The first sentence is Theorem 8.1 evaluated at $m = 4$, $k = 2$, $Q = 1$, together with the arithmetic $3/4 < 5/6$, and is machine-checked; so is the diagonal, where Theorem 8.1 gives $1/(k+1) = 1/m$ at $k = m-1$, which rounds to the printed $0.333, 0.250, 0.200, 0.167, 0.143, 0.125$ for $m = 3, \ldots, 8$. The claim about the 21 cells is *not* machine-checked: it is an external computation, described in Section 12, in which each cell's linear program was generated from the definitions of [1] and solved with `lp_solve` 5.5.2.5, whose eight-digit output was then rationalised by a smallest-denominator approximation. Its results are in Table 1. ☐

Table 1. The cells of Table 6 of [1] (printed), the optimum of the relaxation with Fixed W^* (reproduced externally), and the plain relaxation of Theorem 8.1 at $Q = 1$. Boldface marks the cells discussed in the text; an em dash marks a cell of the plain relaxation that was not part of the numerical reproduction (its value is given by Theorem 8.1 in any case).

m	k	printed	Fixed W^*	plain (2)
3	2	0.333	1/3	1/3
4	2	**0.750**	3/4	**5/6**
4	3	0.250	1/4	1/4
5	2	0.886	31/35	1
5	3	0.667	2/3	3/4
5	4	0.200	1/5	1/5
6	2	0.972	35/36	13/12
6	3	0.833	5/6	11/12
6	4	0.625	5/8	7/10
6	5	0.167	1/6	1/6
7	2	**0.999**	**1**	17/15
7	3	0.901	82/91	1
7	4	**0.777**	**7/9**	13/15
7	5	0.600	3/5	2/3
7	6	0.143	1/7	1/7
8	2	1.000	1	—
8	3	0.951	0.95089286	—
8	4	0.865	0.86548223	—
8	5	0.750	3/4	—
8	6	0.583	7/12	—
8	7	0.125	1/8	—

A few remarks on Table 1. The fractions in the Fixed-W^* column are small-denominator rationalisations of the solver's eight-digit output; the two cells at $m = 8$ with no small denominator are left as decimals rather than asserted as fractions. And the printed cells are the reproduced values *rounded* to three decimals, not truncated: truncation would print 0.885, 0.666, 0.166, 0.142, 0.950 at $(5, 2)$, $(5, 3)$, $(6, 5)$, $(7, 6)$, $(8, 3)$, where the table has 0.886, 0.667, 0.167, 0.143, 0.951. Two cells escape the rule. At $(m, k) = (7, 4)$ the reproduced optimum is $7/9 = 0.\overline{7}$, which rounds to 0.778 where the table prints 0.777 — its truncation, and the only cell printed that way. The one substantive difference is $(m, k) = (7, 2)$, where the exact optimum is 1 and the table prints 0.999 — a solver tolerance in that run, not a disagreement about mathematics.

10. Controls

Every result above is an equality or a refutation about one small linear program, so the failure mode to guard against is a misplaced quantifier that would make the statements true for the wrong reason. The following checks were computed from the same definitions as the results they guard.

Proposition 10.1 (negative controls)**.** *Let m and k be natural numbers.*

(1) Too small refuted. *For $m > k$, no feasible γ keeps every ranking's covered mass at or below $1/(k + 2)$; the optimum of Theorem 4.1 is not smaller than stated.*

(2) Too large refuted. *For $m > k \geq 1$ the uniform γ keeps every ranking's covered mass strictly below $1/k$; the optimum is not larger than stated.*
(3) The hypothesis $k < m$ is load-bearing. *When $m \leq k$ the simplified dual has no feasible point at all, so its optimum is not a number.*
(4) Both optimum notions pin one value. *If v and v' both satisfy Definition 2.2 then $v = v'$, and likewise for the literal reading; so Theorems 4.1 and 5.1 determine their values.*
(5) The witnesses are sharp, and the diagonal is tight. *The elections of Theorem 6.1 do not witness more than they are used for: $E_{4,2}$ does not witness $\alpha^*(4,2) \geq 1/2$, and $E_{6,4}$ does not witness $\alpha^*(6,4) \geq 1/4$. Read at $k = 2$, where $m = k+1$, the cyclic election $E_{3,1}$ witnesses $\alpha^*(3,2) \geq 1/3$ but not $\alpha^*(3,2) \geq 2/5$ — and $1/3 = u^*_{3,2}$, so on the diagonal the failure of Theorem 6.1 does not occur.*
(6) The definitions are not vacuous. *The relation $c \succ_s W$ is neither always true nor always false — under the identity ranking of three candidates, candidate 0 beats $\{1,2\}$ and candidate 2 does not beat $\{0,1\}$ — and $\mathrm{Cov}_{m,k}(s)$ is nonempty whenever $k+1 \leq m$, so Lemma 3.1 is not counting an empty family.*

Item (5) is the one that matters most. At $m = k + 1$ each committee has a single challenger, so the selector equation of (2) forces $y \equiv 1$ and (2) is not a relaxation but a restatement of (1); the Q term of Theorem 8.1 vanishes identically, and the coefficient of Q in the dual objective vanishes with it, since $m - k = 1$ turns the bound of Theorem 7.1 into an equality. On the diagonal, in other words, the elimination that Section 7 refutes is valid, and the simplified dual is the true dual. Theorem 8.1 then reads $\alpha^*(k+1,k) = 1/(k+1) = u^*_{k+1,k}$, so weak duality holds there; at $(m,k) = (3,2)$ item (5) confirms it from the other side, $E_{3,1}$ attaining $1/3$ and not exceeding it. So the failure documented in Section 6 is genuinely an off-diagonal phenomenon and not an artefact of the encoding — which is also what the diagonal of the source's own table, $1/m$ at $k = m-1$, suggests. (Theorem 8.1 is machine-checked; the sentence identifying (2) with (1) at $m = k+1$ is an ordinary argument and is not.)

11. What is not settled

Nothing in this paper bounds the Condorcet dimension. We have evaluated one linear program and examined one derivation. Whether every election admits a Condorcet winning set of size 4 — the question in the title of [1] — remains exactly as open as before: the best published upper bound is 5 [4], the best lower bound is 3 [2], and [5] states the gap as open as of August 2026. In particular:

- Corollary 4.2 does not prove $\alpha \leq 2/k$. It proves the conjectured inequality about u^*_k, and Theorem 6.1 shows that this inequality does not transfer to α^*.
- Corollary 8.2 does not say that the Condorcet dimension is not 4. It says that this particular relaxation, with this big-M constant, cannot decide it, whichever dual certificate one looks for. A different relaxation — a tighter encoding of the selector logic, or a formulation without a big-M constant — is not touched by anything here.
- Nothing here evaluates the multiset variant of the program that [1] also tabulates: there the multiplicity of a candidate in a committee changes the combinatorics, and whether the closed form of Theorem 4.1 survives is open.
- The values $\alpha^*(m,k)$ themselves are computed nowhere in the formal development. Theorem 6.1 gives lower bounds only; the matching upper bounds at those parameters are the MILP values of [1], which were re-solved here only as the external computation of Section 12, and which no theorem of this paper certifies.

12. Provenance

Theorem 4.1, Corollary 4.2 and Theorem 8.1 are the results offered as new. Theorems 5.1, 6.1 and 7.1, Corollary 8.2 and Proposition 9.1 are corrections to [1]; they are new in the sense that no one has written them, but they are consequences of the first three together with the source's own data, and we do not count them as new mathematics. Proposition 10.1 collects controls.

The compute-first check. Before any of the above was proved, the three programs were solved numerically from the definitions of [1], so that the closed forms were confirmed rather than hoped for; the figures below are from a second, independent pass made while this paper was being checked, with `lp_solve` 5.5.2.5 on one workstation core. The simplified dual was solved for all $3 \leq m \leq 7$ and $1 \leq k < m$ — twenty cells — under both readings of its index set: every cell returned $1/(k+1)$ under the reading of Definition 2.2 and 0 under the literal one, which is Theorems 4.1 and 5.1 numerically. The plain relaxation (2) was solved at $Q = 1$ for all $3 \leq m \leq 7$ and $2 \leq k < m$ — fifteen cells — and every cell matched Theorem 8.1 exactly. The Fixed-W^* relaxation was solved for all 21 cells of Table 1.

The MILP of [1] was re-solved at every cell of its Table 3 with $m \leq 6$ — that is, at every cell up to $m = 6$ that the table prints as a number rather than as an interval — giving 1/3 at $(3, 2)$; 2/5 and 1/4 at $m = 4$; 7/15, 3/10 and 1/5 at $m = 5$; and 1/2, 1/3 and 3/13 at $m = 6$. Rounded, these are exactly the printed 0.333; 0.400, 0.250; 0.467, 0.300, 0.200; 0.500, 0.333, 0.231. The three witnesses of Theorem 6.1 with $k \geq 2$ were extracted from the integral solutions at $(4, 2)$, $(5, 3)$ and $(6, 4)$ and then re-verified from scratch against full rankings; the two witnesses at $k = 1$ are printed in [1] and were only transcribed. The $m = 7$ row was not reproduced: [1] itself reports intervals rather than values at $(7, 3)$ and $(7, 4)$, its own solver having timed out there, and the remaining cell $(7, 2)$ — 2520 ranking classes, 21 committees, 105 binary selectors — did not finish inside the budget allowed here. For scale, [1] solves these instances with Gurobi 12.0.3 under the 720 s cap its Table 3 caption names, on cluster nodes with 32 physical cores and 512 GB of RAM. Nothing in this paper depends on the $m = 7$ cells.

All of these computations are external to the formal development and are not proofs; they are reported because they are what makes the theorems' statements the right ones to have proved. The ranking enumeration used the "reducing permutations" optimisation of [1], keying rankings by their top $m - k$ prefix, which that paper shows is exact for committees of size k.

Literature. Searched on 2026-09-02 and re-checked on 2026-09-03. [1] is a workshop paper of April 2026 with, at both dates, no recorded citations: OpenAlex has no record of the identifier 2604.19851 at all, and the sole OpenAlex citer of [3] is an unrelated survey of Jérôme Lang's work. In a local full-text corpus of arXiv preprints with coverage through the 2026-08-28 announcement, 22 papers contain the phrase "Condorcet winning set" and 6 contain "Condorcet dimension"; every one of them posted after [1] was extracted and searched for the strings `2604.19851`, "Zilberstein", "simplified dual", "dual linear program", "Dual Bound", "Is Four Enough", "weak duality" and "LP relaxation", with no hit anywhere. A live arXiv metadata query returns exactly nine papers for `all:"Condorcet winning set"` in the whole of arXiv and four for `all:"Condorcet dimension"`; the three posted after [1] were read and none mentions it, its conjecture, or LP duality for this problem. Web search for the closed form $1/(k+1)$ as the optimum of a program of this shape returned only [1] and the known-results papers. A zbMATH query for "Condorcet winning set" returns 16 records, none of them [1] — a workshop paper without a DOI is in none of these indexes — and the four carrying a 2026 date are [4], an exposition of it (arXiv:2606.14666, read above), and the journal versions of arXiv:2410.09201 and arXiv:2601.10506, both posted before [1] was. As Remark 4.3 says, the constant itself is classical and we claim no priority in it; what is claimed is its optimality for this program, and in particular the matching lower bound.

The channels not searched, which bounds the claim: MathSciNet, Google Scholar, and any publication absent from arXiv, from zbMATH and from the corpus above. "Not found" is not "new", and what most weakens the claim is simply that the program is four months old, so a literature about it has had no time to exist.

One attribution note. [1] credits [3] with the lower bound: "Generalizing this, [Charikar et al.] show that if a size-k committee is always guaranteed, then α must be at least $\frac{2}{k+1}$", and later calls it "the lower bound established by" that paper. [3] does carry the bound — in its Section 1.3, "the general lower bound that α-undominated committees of size k do not always exist when $\alpha < \frac{2}{k+1}$", and again as a theorem with a full proof in its Appendix B — but in both places it credits the statement to a 2014 MathOverflow post of Sam Zbarsky, `mathoverflow.net/q/157982`, rather than claiming it. Quotations from [3] are from v4 of its e-print, read 2026-09-03. We record the chain because a reader tracing the constant will meet it; nothing in this paper rests on it.

13. Verification

Every theorem, corollary and proposition above — with the explicitly flagged exceptions, Remark 6.2, the 21-cell reproduction of Proposition 9.1, and the sentence in Section 10 that identifies (2) with (1) on the diagonal $m = k+1$, each of which says at the point where it stands that it is outside the formal development — has been machine-checked in Lean 4 with Mathlib, in a repository that is private at the time of writing (repository reference to be added at submission). The development is four modules carrying 35 headline declarations, and contains no `sorry` and no declared axiom; each declaration carries the axiom set {`propext`, `Classical.choice`, `Quot.sound`}, except one control that uses only {`propext`, `Quot.sound`}. Compiled evaluation is not used anywhere: Lean's `native_decide` appears in no module, so no statement rests on trusting a compiler. The only exhaustive computations inside the development are the finite checks of Theorem 6.1 and of Proposition 10.1 (5), together with the three-candidate check in item (6): for each of the five elections, the assertion is decided by the kernel over all $\binom{m}{k} \leq 15$ committees of the relevant size, with the challenger quantifier over at most six candidates and the voter counts summed over at most thirteen listed rankings. Elections are carried with integer voter weights and every witness statement has the shape $a \cdot (\text{voters}) \leq b \cdot (\text{support})$ over $\mathbb{N}$, so no rational comparison is ever evaluated by the kernel and no set of rankings is ever enumerated; the enumerations the kernel does perform run over the subsets of a six-element candidate set and over that set itself. Everything else — Lemmas 3.1 and 3.2, Theorems 4.1, 5.1, 7.1, 8.1, Corollaries 4.2 and 8.2, and the general items of Proposition 10.1 — is proved uniformly in m, k and Q. The cost is negligible: 47.9 seconds of processor time for the whole development in import order, with peak resident set equal to the baseline for importing the library alone. It builds under the Lean toolchain `leanprover/lean4:v4.32.0` with Mathlib at its `v4.32.0` tag [7].

References

[1] I. Zilberstein, R. E. Berker, G. Li and R. Martins, *Is Four Enough? Automated Reasoning Approaches and Dual Bounds for Condorcet Dimensions of Elections*, arXiv:2604.19851v1 (submitted 21 April 2026); GAIW-26, the 8th Games, Agents, and Incentives Workshop, held with the workshops of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2026), Paphos. Primary arXiv category `cs.GT`, cross-listed `cs.AI`. All quotations here are from version 1, read from the LaTeX e-print; version 1 is the only version and the abstract page carries no journal reference (checked 2026-09-02 and again 2026-09-03). Equation, table, figure and conjecture numbers cited here are those of the compiled v1 document, which runs to nine pages.

[2] E. Elkind, J. Lang and A. Saffidine, *Condorcet winning sets*, Social Choice and Welfare **44** (2015), no. 3, 493–517, `doi:10.1007/s00355-014-0853-4`; Zbl 1318.91066. Bibliographic data confirmed against Crossref, dblp and zbMATH on 2026-09-03. The paper itself was not consulted here: it is cited for the definition of a Condorcet winning set and for the lower bound 3, and both attributions follow [1] and the abstract of [3], which records that Elkind, Lang and Saffidine "showed that winning sets of size 2 may not exist, but sets of size logarithmic in the number of candidates always do".

[3] M. Charikar, A. Lassota, P. Ramakrishnan, A. Vetta and K. Wang, *Six Candidates Suffice to Win a Voter Majority*, in: Proceedings of the 57th Annual ACM Symposium on Theory of Computing (STOC '25), ACM, 2025, pp. 1590–1601, `doi:10.1145/3717823.3718235`; arXiv:2411.03390v4. Page range and author list confirmed against Crossref on 2026-09-03; the Section 1.3 and Appendix B passages quoted in Section 12 were read from the v4 e-print the same day.

[4] H. Song, T. Nguyen and Y.-S. Lin, *A Few Good Choices*, in: Proceedings of the 2026 Annual ACM–SIAM Symposium on Discrete Algorithms (SODA), SIAM, 2026, pp. 4861–4874, `doi:10.1137/1.9781611978971.175`; preprint arXiv:2506.22133v2, whose arXiv metadata lists the authors in the order T. Nguyen, H. Song, Y.-S. Lin. The current best upper bound, $k = 5$, for the size of a Condorcet winning set. Confirmed against Crossref and dblp on 2026-09-03.

[5] M. Charikar, P. Ramakrishnan and K. Wang, *Pessimal Elections for Approximately Dominating Sets*, preprint, arXiv:2608.06872v1 (submitted 7 August 2026). Quoted in Section 1, from a footnote of its introduction, for the standing of the size-4 question as of August 2026. Author list, posting date and the quoted footnote were all read on 2026-09-03, the first two from the arXiv metadata and the last from the e-print.

[6] Y. Lin, S. Qin, K. Wang and L. Xia, *Winning in the Limit: Average-Case Committee Selection with Many Candidates*, arXiv:2602.04815v1 (submitted 4 February 2026). Cited for its related-work table, which records the large-voter impartial-culture threshold $1 - 1/(k+1)$ for winning committees and credits it to [2], Corollary 2. Author list and date read from the arXiv metadata on 2026-09-03.

[7] 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, pp. 367–381, `doi:10.1145/3372885.3373824`.