THE TOP TWO SCORES OF A ROUND-ROBIN TOURNAMENT: EXACT DISTRIBUTIONS, AND WHEN THE MAXIMUM WINNER'S MARGIN IS ATTAINED

KOREA SUPERINTELLIGENCE LABS

Abstract. In a round-robin tournament of n equally strong players, let $s_{(n)}$ and $s_{(n-1)}$ be the largest and the second largest score. Both are named as objects of study in the recent literature, which calls the exact distribution computationally intractable even for small n and derives asymptotics instead. It is not intractable: carrying only the sorted multiset of scores gives a dynamic program whose state count is the number of realisable score sequences, and we compute the joint law together with the marginals of the runner-up score, of the winner's margin $s_{(n)} - s_{(n-1)}$, and of the number of co-champions, exactly, for the classical model up to $n = 13$, the chess model up to $n = 9$ and the soccer model (3 points for a win, 1 each for a draw) up to $n = 7$. Two consequences are uniform in n. First, the winner is unique exactly when the margin is positive, so the margin distribution's mass at 0 is the complement of the unique-winner count in every outcome model. Second, our main theorem, a characterisation of the largest margin any tournament realises. Write M for the most one player can take from a single match and $\mu(k)$ for the smallest winner's score over all k-player tournaments of the model, the *minimax score*. On $n \geq 2$ players no margin exceeds $M(n-1) - \mu(n-1)$, with no hypothesis on the model whatever; and that bound is the largest realised margin *if and only if* the model contains a total-domination result, one awarding M points to one player and nothing to the other. If moreover some pair of results splits the smallest match total m as evenly as possible in both orders, then $\mu(k) = \lceil m(k-1)/2 \rceil$ and the bound becomes the closed form $M(n-1) - \lceil m(n-2)/2 \rceil$, which is $\lfloor kn/2 \rfloor$ in the model where a match is worth k points and every split of it is possible — hence $\lfloor n/2 \rfloor$ classically and n in the chess model — and $2n - 1$ in the soccer model. Without that extra hypothesis the closed form can be strictly too large: when a win is worth $g \geq 2$ and draws are impossible, the largest margin is $g\lfloor n/2 \rfloor$ at every $n \geq 2$. Everything asserted here is machine-checked in Lean 4, the margin results from the definitions with no computation anywhere in their proofs.

1. Introduction

The objects. Fix $n \geq 2$ players, each of whom meets each other player exactly once. A match is settled by one *outcome* drawn from a finite list $\mathcal{O}$ of pairs of non-negative integers: the outcome (a, b) awards a points to the first-named player of the match and b to the other. Three lists occur throughout:

$$D_k = \{(a, k-a) : 0 \leq a \leq k\}, \qquad \mathcal{O}_{\text{soc}} = \{(3,0), (1,1), (0,3)\}.$$

D_1 is the *classical* model, one point to the winner and nothing to the loser; D_2 is the *chess* model, win/draw/loss scored $1, \frac{1}{2}, 0$ with every score doubled so that it stays an integer; $\mathcal{O}_{\text{soc}}$ is the *soccer* model, three points for a win and one each for a draw. The family D_k is the model $D_k = \{0, 1/k, \ldots, 1\}$ of Malinovsky [13, §1] after clearing denominators; there the two players' scores in a match are required to sum to 1, which in integer scores reads $a + b = k$. The soccer model violates that requirement and is included precisely because it does.

A *tournament* is a choice of one outcome for each of the $\binom{n}{2}$ matches, so the sample space has $|\mathcal{O}|^{\binom{n}{2}}$ elements; *equally strong players* means the uniform distribution on it. The *score sequence* is the vector $(s_1, \ldots, s_n)$ of the players' point totals, and $s_{(1)} \leq \cdots \leq s_{(n)}$ are its order statistics. This paper is about the top of that vector: the winner's score $s_{(n)}$, the runner-up's score $s_{(n-1)}$, the *margin* $s_{(n)} - s_{(n-1)}$, and the number of *co-champions* $\#\{i : s_i = s_{(n)}\}$.

Where the questions come from. MacMahon [10, 11] computed the exact distribution of the score sequence by hand a century ago — classically for $n \leq 9$, in the chess model for $n \leq 6$ — by the calculus of symmetric functions. Malinovsky [13] takes those two ranges as the state of the exact art and writes in §1 that

> "even for small values of n, determining the exact distribution is in general computationally infeasible due to the astronomical size of the sample space, which, for example, in the classical round-robin tournament equals $2^{\binom{n}{2}}$",

and in its abstract that

> "[s]ince the exact distribution is computationally intractable even for small values of n, we derive asymptotic results as the number of players n tends to infinity".

The paper's own content is a Poisson approximation for the number of players whose normalised score exceeds a threshold, from which its Corollary 2 deduces, for every fixed $j \geq 0$, a limit law for the j-th order statistic from the top. Nothing there is exact and nothing there is finite n. The same author's earlier work on the winner's score [14] states twice that the exact distribution for general n seems impossible to obtain and tabulates Monte-Carlo estimates instead; its subsection on the second and third largest scores gives moments of a Gumbel limit. A follow-up [12] proves that the top $k(n)$ scores are distinct with probability tending to 1 whenever $k(n) \to \infty$ slowly enough; in particular the top two scores are distinct asymptotically almost surely, i.e. the margin is positive with probability tending to 1.

The infeasibility claim is false, and it was already false: the sample space is a red herring, because the score multiset of the players already seated is a sufficient state for adding one more player. Zeilberger, writing as Shalosh B. Ekhad [3], ran essentially that dynamic program in Maple in 2016 and published the full classical table to $n = 12$, the full chess table to $n = 9$ and the full soccer table to $n = 7$. What was missing was not the tables but everything read *off* the top of them: his package computes the table and its mode, and no order statistic other than the maximum.

What is settled here.

(1) The exact joint law of $(s_{(n)}, s_{(n-1)})$ and the three marginals — runner-up score, margin, number of co-champions — for the classical model at $n \leq 13$, the chess model at $n \leq 9$ and the soccer model at $n \leq 7$ (Section 4). These are counts over the whole sample space, not properties of an algorithm: the dynamic program is proved to compute the enumeration it replaces (Section 2), so no statement in this paper mentions the algorithm. Every one of these numbers is a finite summation over a table that [3] already published, and we claim for them exactness and verification, not novelty.

(2) Two facts uniform in n and in the outcome model (Section 3): on any score vector of at least two players the winner is unique if and only if the margin is positive, and consequently

$$\#\{\text{tournaments with margin } 0\} + \#\{\text{tournaments with a unique winner}\} = |\mathcal{O}|^{\binom{n}{2}}$$

at every $n \geq 2$. The second summand is a published sequence in the classical model, known far past any range computed here, so the margin distribution is pinned at 0 for those n as well.

(3) The maximum winner's margin (Section 5), which is where the new results are. Write $M = M(\mathcal{O})$ for the largest number of points one player can take from a single match, and let the *minimax score* $\mu_{\mathcal{O}}(k)$ be the smallest winner's score over all k-player tournaments of the model. Then at every $n \geq 2$ no tournament on n players has margin greater than $M(n-1) - \mu_{\mathcal{O}}(n-1)$, with no hypothesis on $\mathcal{O}$ at all; and that bound is the greatest margin actually realised *if and only if* $(M, 0) \in \mathcal{O}$ or $(0, M) \in \mathcal{O}$, i.e. exactly when the model contains a result in which one player takes everything and the other nothing. Both disjuncts are needed: there are models attaining the bound through each one alone. The characterisation also settles, negatively and at every n at once, models where the bound could previously be refuted only one n at a time. We could not find either half in the literature, in any model; Section 5 records the searches. Every proof there is elementary and none uses any computation.

(4) The value of that bound, and the price of stating it as a closed form (Section 5 again): $\mu_{\mathcal{O}}(k) = \lceil m(k-1)/2 \rceil$ as soon as some pair of results splits $m = m(\mathcal{O})$ as evenly as possible in both orders, giving $\lfloor kn/2 \rfloor$ in D_k for every k and $2n-1$ in the soccer model; and, without that hypothesis, models where the closed form overshoots at every odd n.
(5) The soccer model as a test of model-genericity (Section 7): the correctness proof of the dynamic program never uses a constant match total, so a model outside the family D_k costs nothing, and its tables reproduce 41 343 published cells exactly.

Section 6 collects the checks — refutations of too-large and too-small claims, a direct enumeration of the sample space, and a comparison against the one published quantity that constrains the co-champion distribution beyond its first column. Section 8 summarises the reproduction of the classical published material, including MacMahon's nine-player erratum. Section 9 describes the formal development.

Changes in this version. Section 5 is reorganised around the minimax score, which the previous version of this paper did not contain; no statement of that section is withdrawn, each being either kept or restated in a form that implies it. What was the main theorem — the largest margin is $M(n-1) - \lceil m(n-2)/2 \rceil$ under two hypotheses (i) and (ii) — is now Corollary 5.10, a consequence of Theorem 5.6 together with the evaluation $\mu_{\mathcal{O}}(k) = \lceil m(k-1)/2 \rceil$ that (ii) supplies (Proposition 5.9). New here: the sharp upper bound, which needs no hypothesis (Proposition 5.3); attainment under (i) alone, which reveals (ii) as a statement about $\mu_{\mathcal{O}}$ and not about margins (Theorem 5.6); the converse, which makes attainment an exact characterisation (Theorem 5.7); the no-draw family, on which the earlier closed form is strictly larger than the truth (Proposition 5.14); and the controls that go with these, among them two models in which the bound is provably not attained at any n (Proposition 5.16), where the previous version could refute attainment only one n at a time, and two one-outcome families that attain it through a single orientation (Proposition 5.15). The question the previous version left open — when exactly the bound is attained — is answered by Theorem 5.7 for the sharp bound $M(n-1) - \mu_{\mathcal{O}}(n-1)$; for the closed form $B(M, m, n)$, the bound of the previous version, the answer is Theorem 5.7 together with the evaluation of $\mu_{\mathcal{O}}$, which Proposition 5.9(2) supplies under (ii) and Proposition 5.14 shows can fail without it. What remains open is Remark 5.8.

2. Preliminaries: tournaments, the score multiset, and the transfer principle

Definition 2.1. An *outcome model* is a finite list $\mathcal{O}$ of pairs $(a, b) \in \mathbb{N}^2$. For $n \geq 0$ let $T_n(\mathcal{O})$ be the set of *tournaments*: functions assigning to each pair $i < j$ of players an element of $\mathcal{O}$, read as "player i takes a points and player j takes b points from that match". Then $|T_n(\mathcal{O})| = |\mathcal{O}|^{\binom{n}{2}}$. The *score* $s_i(t)$ of player i in $t \in T_n(\mathcal{O})$ is the sum of the points t awards to i, and the *score multiset* is $\sigma(t) = \{\{s_1(t), \dots, s_n(t)\}\}$.

For a multiset μ of n non-negative integers we write $\mathrm{dist}_{\mathcal{O}}(n, \mu)$ for the number of $t \in T_n(\mathcal{O})$ with $\sigma(t) = \mu$; this is the *exact distribution of score sequences* of [10, 13], the probability of μ being $\mathrm{dist}_{\mathcal{O}}(n, \mu)/|\mathcal{O}|^{\binom{n}{2}}$. We also write, for the two model constants used throughout,

$$M(\mathcal{O}) = \max_{(a,b) \in \mathcal{O}} \max(a, b), \qquad m(\mathcal{O}) = \min_{(a,b) \in \mathcal{O}} (a + b),$$

so that $(M, m) = (1, 1)$ classically, $(2, 2)$ in the chess model and $(3, 2)$ in the soccer model.

2.1. **The dynamic program.** Seat the players one at a time. If the players already seated have score multiset μ and a newcomer plays them all, the resulting multiset of extended score multisets depends on μ only — not on which player had which score, and not on the tournament that produced μ. Carrying the sorted score multiset as the state therefore suffices, and the number of states at stage n is the number of realisable score sequences: 1, 1, 1, 2, 4, 9, 22, 59, 167, 490, 1486, 4639, 14 805, 48 107 for classical $n = 0, \dots, 13$ (sequence A000571 of [17]), against $2^{\binom{13}{2}} = 2^{78}$ tournaments. We write $\mathcal{T}(\mathcal{O}, n)$ for the resulting weighted list of sorted score vectors.

Proposition 2.2 (correctness of the dynamic program)**.** *For every outcome model $\mathcal{O}$ and every n, the weighted list $\mathcal{T}(\mathcal{O}, n)$ represents exactly the multiset $\{\{\sigma(t) : t \in T_n(\mathcal{O})\}\}$: for every multiset μ, the weight $\mathcal{T}(\mathcal{O}, n)$ carries at μ is* $\mathrm{dist}_{\mathcal{O}}(n, \mu)$.

Proof sketch. Induction on n, with no computation anywhere. The one substantial step is the permutation lemma: if v and w are score lists that agree as multisets, the weighted lists of one-player extensions of v and of w agree. It is proved by induction over the permutation relation, the transposition case being commutativity of multiset insertion together with commutativity of addition. Two bookkeeping facts complete the induction: the enumeration of all rows of a newcomer against j seated players matches the extension operator, and the sorting-and-merging step that makes the program fast is invisible to the represented multiset, so no part of the argument depends on how that step is implemented. □

The consequence we use everywhere is that any statistic which cannot see the players' names is computable from the table.

Proposition 2.3 (transfer principle)**.** *Let P be a predicate on score lists that is invariant under permutation. Then for every $\mathcal{O}$ and n,*

$$\#\{t \in T_n(\mathcal{O}) : P(s_1(t), \ldots, s_n(t))\} \;=\; \sum_{(v,c)\in\mathcal{T}(\mathcal{O},n),\ P(v)} c.$$

In particular, taking P identically true, the total weight of $\mathcal{T}(\mathcal{O}, n)$ is $|\mathcal{O}|^{\binom{n}{2}}$ at every n — an identity, not a numerical check.

Three instances of Proposition 2.3 are immediate: the count at one score vector, the winner's-score marginal $\mathrm{maxDist}_{\mathcal{O}}(n, b) = \#\{t : s_{(n)}(t) = b\}$, and the unique-winner count. The point of Section 3 is that the top two order statistics are three further instances, obtained at negligible cost.

Remark 2.4. The unique-winner event is the subject of a conjecture of Epstein, stated without proof in the 1967 first edition of his book [4, Ch. 9] and discussed in [13, §1]: the probability that a classical round-robin tournament of equally strong players has a unique winner tends to 1 as $n \to \infty$. It was proved by Malinovsky and Moon [15], whose abstract and §1 record the attribution and the history. Amir and Malinovsky [1] carry the same limit past the classical model: if the two players' scores in a match sum to 1 and take their values in a countable set $D \subseteq [0, 1]$ with $|D| > 1$, every value having positive probability, then in a tournament of equally strong players the probability of a unique maximum score again tends to 1. Their setting contains each D_k, after dividing scores by k; it does not contain the soccer model, in which the two players' scores do not sum to a constant (Proposition 7.1). The exact counts are sequence A013976 of [17], known to $n = 20$. Both of those papers concern the uniqueness of the maximum score; neither says anything about the margin by which it is attained, which is the subject of Section 5.

3. The top two order statistics

Definition 3.1. For a list v of scores put

$$\max(v) = s_{(n)}, \qquad \mathrm{second}(v) = \max\bigl(v \setminus \{\text{one copy of } \max(v)\}\bigr) = s_{(n-1)},$$

$$\mathrm{gap}(v) = s_{(n)} - s_{(n-1)}, \qquad \mathrm{champ}(v) = \#\{i : v_i = \max(v)\}.$$

The runner-up is taken *with multiplicity*: on $(3, 3, 1)$ it is 3, not 1, so a tie at the top gives margin 0 and two co-champions.

Proposition 3.2. *Each of* second*,* gap*,* champ *and the pair* $(\max, \mathrm{second})$ *is invariant under permutation of the score list, hence Proposition 2.3 applies to each. Consequently the four counts*

$$\#\{t : s_{(n-1)}(t) = b\}, \quad \#\{t : \mathrm{gap}(t) = d\}, \quad \#\{t : \mathrm{champ}(t) = k\}, \quad \#\{t : s_{(n)}(t) = a,\ s_{(n-1)}(t) = b\}$$

over all $|\mathcal{O}|^{\binom{n}{2}}$ *tournaments are read off* $\mathcal{T}(\mathcal{O}, n)$*. Moreover* $\operatorname{second}(v) \le \max(v)$*, so the margin is a genuine difference, and the co-champion distribution refines the unique-winner count: its value at* $k = 1$ *is the number of tournaments with a unique winner, at every* n *and in every model.*

Proof sketch. Permutation invariance of max is standard; for second it follows because deleting one copy of a common element from two lists that agree as multisets again leaves lists that agree as multisets, and for champ because counting occurrences of a fixed value is invariant. The inequality $\operatorname{second}(v) \le \max(v)$ holds since every element of the reduced list is an element of v. The last claim is a definitional identity: "the maximum is attained exactly once" *is* "champ = 1". ☐

The one statement in this section with content is the following, and it is the reason the margin and the classical unique-winner sequence determine each other.

Theorem 3.3. *Let* v *be a score list of length at least* 2*. Then the maximum of* v *is attained exactly once if and only if* $\operatorname{gap}(v) > 0$.

The length hypothesis cannot be dropped: on the one-player list (0) the maximum is attained once while the margin is 0 (Proposition 6.1(2)).

Proof sketch. Both directions rest on two facts: a non-empty list attains its maximum, and deleting one occurrence of a value drops its multiplicity by exactly one. If $\max(v)$ occurs once, then after the deletion it does not occur at all; the reduced list is non-empty because v has at least two entries, so it attains its own maximum at some entry which is $\le \max(v)$ and $\ne \max(v)$, whence $\operatorname{second}(v) < \max(v)$. If $\max(v)$ occurs at least twice, it still occurs after the deletion, so $\operatorname{second}(v) \ge \max(v)$ and the margin is 0. ☐

Corollary 3.4. *For every outcome model* $\mathcal{O}$ *and every* $n \ge 2$,

$$\#\{t \in T_n(\mathcal{O}) : \operatorname{gap}(t) = 0\} \;+\; \#\{t \in T_n(\mathcal{O}) : \textit{the winner is unique}\} \;=\; |\mathcal{O}|^{\binom{n}{2}}.$$

Proof sketch. Split $T_n(\mathcal{O})$ by the unique-winner predicate and apply Theorem 3.3 to each tournament, the length hypothesis being available because a tournament on n players has n scores. The count of the complement is then the count of margin 0. ☐

Corollary 3.4 is uniform in n, so in the classical model it determines the margin distribution at $d = 0$ for every $n \le 20$, from A013976 — seven values beyond the largest table below.

4. The exact distributions

Theorem 4.1. *The joint distribution of* $(s_{(n)}, s_{(n-1)})$ *over all labelled tournaments, and with it the marginals of the runner-up score, of the margin and of the number of co-champions, are known exactly for the classical model at* $2 \le n \le 13$ *and the chess model at* $2 \le n \le 9$*. The row sums of the joint matrix are the winner's-score marginal and its column sums are the runner-up marginal. Tables 1 and 2 give two of the marginals in the classical model; the full data are in the formal development. The same four statistics in the soccer model,* $2 \le n \le 7$*, are Proposition 7.1.*

Proof sketch. For each model and each n one evaluation of the dynamic program produces $\mathcal{T}(\mathcal{O}, n)$, and the four statistics are read off it by filtering that finite weighted list; the evaluation is a computation inside the proof assistant, delegated to compiled code, and Propositions 2.2 and 3.2 convert the resulting statement about $\mathcal{T}(\mathcal{O}, n)$ into the statement about the $|\mathcal{O}|^{\binom{n}{2}}$ tournaments. The exhaustive part is exactly that evaluation: 48 107 states at classical $n = 13$ and 26 376 at chess $n = 9$, together with w^2 further passes over the state list for the joint matrix, w being the number of attainable scores (n classically, $2n - 1$ in the chess model). No claim in this theorem is asymptotic and none is sampled. ☐

n	$d=0$	$d=1$	$d=2$	$d=3$	$d=4$
2	0	2	0	0	0
3	2	6	0	0	0
4	32	24	8	0	0
5	424	440	160	0	0
6	12 224	16 320	4 080	144	0
7	878 928	941 808	240 128	36 288	0
8	108 194 304	118 304 256	35 976 192	5 939 584	21 120

Table 1. The classical margin marginal: the number of tournaments on n players in which the winner's margin is exactly d. Each row sums to $2^{\binom{n}{2}}$. The last non-zero entry of row n sits at $d = \lfloor n/2 \rfloor$, which is Theorem 5.17; the mode moves from $d = 0$ to $d = 1$ at $n = 5$ and stays there.

n	$k=1$	$k=2$	$k=3$	$k=4$	$k=5$	$k=6$	$k=7$
2	2	0					
3	6	0	2				
4	32	24	8	0			
5	600	360	40	0	24		
6	20 544	6 720	2 960	2 400	144	0	
7	1 218 224	507 360	305 760	62 160	1 008	0	2 640

Table 2. The classical co-champion distribution: the number of tournaments on n players whose maximum score is attained by exactly k of them. The column $k = 1$ is A013976. The entry at $k = n$ counts regular tournaments (A007079) and vanishes for even n, so the distribution has interior zeros and is not monotone in k.

At the top of the classical range the margin marginal on twelve players, over all 2^{66} tournaments, is

$$\begin{aligned} d=0:&\ 27\,827\,016\,988\,869\,062\,656, & d=4:&\ 335\,352\,455\,668\,015\,104,\\ d=1:&\ 32\,445\,472\,803\,775\,156\,224, & d=5:&\ 21\,613\,585\,606\,609\,920,\\ d=2:&\ 10\,725\,370\,013\,860\,896\,768, & d=6:&\ 579\,018\,101\,760,\\ d=3:&\ 2\,432\,149\,868\,040\,364\,032, & d\geq 7:&\ 0. \end{aligned}$$

Remark 4.2 (what is and is not new in Theorem 4.1). Every number above is determined by the score-sequence tables of [3], which are published in full for classical $n \leq 12$, chess $n \leq 9$ and soccer $n \leq 7$: each of our four statistics is a function of the score multiset, so summing the published table's cells over the level sets of that function reproduces it. Classical $n = 13$ is one step past the published range, and it is still a re-run of a published algorithm. What we add is that the values are exact, that they are stated about the sample space rather than about an algorithm, and that both the values and the reduction are machine-checked. We searched the On-Line Encyclopedia of Integer Sequences [17] for all four statistics — nineteen numerical queries and six name searches, with four positive controls that all matched — and found none of them; that is a statement about the encyclopedia, not a novelty claim.

5. The maximum winner's margin

The margin marginals of Theorem 4.1 stop dead, in every model, at a value with a uniform description. The description has two layers, and the first result of this section is that they should be separated. Write

$$B(M, m, n) \;=\; M(n-1) - \left\lceil \frac{m(n-2)}{2} \right\rceil,$$

which is the closed form the computed tables suggest, and introduce the quantity that turns out to carry the actual content.

Definition 5.1 (the minimax score)**.** For an outcome model $\mathcal{O}$ and $k \geq 0$ put

$$\mu_{\mathcal{O}}(k) \;=\; \min_{t \in T_k(\mathcal{O})} \; s_{(k)}(t),$$

the smallest winner's score over all k-player tournaments of the model, with $\mu_{\mathcal{O}}(k) = 0$ when $T_k(\mathcal{O})$ is empty. It is a minimum over a sample space rather than a formula, it does not mention margins, and it is not $m(\mathcal{O})$.

Classically $\mu_{D_1}(k) = \lceil (k-1)/2 \rceil$ is the winner's score in a regular or near-regular tournament, so the minimax score is a familiar object in the classical model; the point of what follows is that it is the right object in *every* model. Throughout this section $M(n-1) - \mu_{\mathcal{O}}(n-1)$ is a difference of natural numbers, and it is non-negative: the winner of an $(n-1)$-player tournament takes at most M from each of its $n-2$ matches, so $\mu_{\mathcal{O}}(n-1) \leq M(n-2) \leq M(n-1)$.

5.1. **The sharp upper bound.**

Lemma 5.2 (deleting the champion)**.** *Let $t \in T_n(\mathcal{O})$ and let i be any player. Deleting i leaves a tournament of the same model on $n-1$ players in which every remaining player's score is at most what it was in t. Consequently, for $n \geq 2$,*

$$s_{(n-1)}(t) \;\geq\; \mu_{\mathcal{O}}(n-1) :$$

in every tournament the runner-up scores at least the minimax score one level down.

Proof sketch. Mathematically this is two sentences: delete the champion, and what is left is an $(n-1)$-player tournament of the same model whose own winner scores no more than the original runner-up, hence at least $\mu_{\mathcal{O}}(n-1)$. Formally the work is in deleting an *arbitrary* player, since the recursion generating $T_n(\mathcal{O})$ only ever removes the most recently seated one; two inductions do it, and the score comparison must be made pointwise, position by position, because the amounts the deleted player handed out differ from position to position. No computation, no restriction on n. ☐

Proposition 5.3 (the upper bound, sharp)**.** *Let $\mathcal{O}$ be any outcome model and let M be any number with $a \leq M$ and $b \leq M$ for every $(a, b) \in \mathcal{O}$. Then for every $n \geq 2$ and every $t \in T_n(\mathcal{O})$,*

$$s_{(n)}(t) - s_{(n-1)}(t) \;\leq\; M(n-1) - \mu_{\mathcal{O}}(n-1),$$

and consequently $\#\{t \in T_n(\mathcal{O}) : \mathrm{gap}(t) = d\} = 0$ for every $d > M(n-1) - \mu_{\mathcal{O}}(n-1)$. There is no other hypothesis: no lower bound on match totals, no assumption about which results $\mathcal{O}$ contains.

Proof sketch. Combine $s_{(n)}(t) \leq M(n-1)$, which is a one-line induction saying that a player has $n-1$ matches and takes at most M from each, with Lemma 5.2. No computation, no restriction on n. ☐

5.2. **Attainment, and the characterisation.**

Attainment is where hypotheses enter, and the content of this section is that exactly one is needed. Call a result *total domination* if it awards M points to one player and nothing to the other, i.e. if it is $(M, 0)$ or $(0, M)$.

Proposition 5.4 (attainment from a total-domination result)**.** *Let $\mathcal{O}$ and M be as in Proposition 5.3 and let $n \geq 2$.*

1. *If $(M, 0) \in \mathcal{O}$, some tournament in $T_n(\mathcal{O})$ has margin exactly $M(n-1) - \mu_{\mathcal{O}}(n-1)$.*
2. *If $(0, M) \in \mathcal{O}$, the same conclusion holds.*

Proof sketch. Both witnesses are two lines and neither involves any computation. Take an $(n-1)$-player tournament t_0 attaining the minimax score — one exists, $T_{n-1}(\mathcal{O})$ being finite and non-empty — and add one player who wins each of his $n-1$ matches by the total-domination result. He scores $M(n-1)$ and hands out nothing, so the others' scores are unchanged, the runner-up scores $\mu_{\mathcal{O}}(n-1)$, and the margin is $M(n-1) - \mu_{\mathcal{O}}(n-1)$.

The two cases differ only in where the new player must be seated, and that is not a formality. A result is an ordered pair whose first coordinate goes to the more recently seated player of the match, so a champion seated *last* needs $(M, 0)$ in each of his matches and one seated *first* needs $(0, M)$; the second construction therefore appends one match to every row of t_0 rather than adding a row, and its score list is again computed in closed form. □

The converse is the new content, and it needs a decomposition of a player's score that the arguments above avoid.

Lemma 5.5 (a player's own matches)**.** *Let $t \in T_n(\mathcal{O})$ and let i be a player. Write $R_i(t)$ for the list of i's $n-1$ matches, each recorded* from i's side *as the pair (points to i, points to the opponent), the opponents listed in the order in which they appear in the score list of t with i's entry removed. Then:*

1. *every entry of $R_i(t)$ either lies in $\mathcal{O}$ or has its coordinate swap in $\mathcal{O}$, and the disjunction cannot be dropped;*
2. *$s_i(t)$ is the sum of the first coordinates of $R_i(t)$;*
3. *deleting i (Lemma 5.2) lowers the other players' scores by* exactly *the second coordinates of $R_i(t)$: the score list of t with i's entry removed is the entrywise sum of the score list of t with i deleted and the list of second coordinates of $R_i(t)$.*

Proof sketch. Three inductions on the seating recursion; no computation. Item (1) is a case distinction: a match against an earlier-seated opponent is recorded in i's own row and is already oriented from i's side, while one against a later-seated opponent is recorded in that opponent's row and is therefore reversed. Items (2) and (3) unfold the definition of the score, (3) being the exact version of the inequality in Lemma 5.2. The design choice worth recording is that both directions of a match are kept in a single entry, so no argument has to certify that two separately indexed quantities came from the same match. □

Theorem 5.6 (the largest margin, sharply)**.** *Let $\mathcal{O}$ be an outcome model with $a \leq M$ and $b \leq M$ for every $(a, b) \in \mathcal{O}$, and suppose $(M, 0) \in \mathcal{O}$. Then for every $n \geq 2$ the set of winner's margins realised by tournaments in $T_n(\mathcal{O})$ has a greatest element, and that element is*

$$M(n-1) - \mu_{\mathcal{O}}(n-1).$$

Proof sketch. Proposition 5.3 and Proposition 5.4(1). □

Theorem 5.7 (attainment characterised)**.** *Let $\mathcal{O}$ be an outcome model with $a \leq M$ and $b \leq M$ for every $(a, b) \in \mathcal{O}$, and let $n \geq 2$. Then the following are equivalent:*

1. *some tournament in $T_n(\mathcal{O})$ has margin exactly $M(n-1) - \mu_{\mathcal{O}}(n-1)$;*
2. *the set of margins realised by tournaments in $T_n(\mathcal{O})$ has greatest element $M(n-1) - \mu_{\mathcal{O}}(n-1)$;*
3. *$(M, 0) \in \mathcal{O}$ or $(0, M) \in \mathcal{O}$.*

In particular the upper bound of Proposition 5.3, which holds in every model, is the true maximum exactly in the models containing a total-domination result; in any other non-empty model it is strictly larger than the largest realised margin, at every $n \geq 2$.

Proof sketch. That (3) implies (1) is Proposition 5.4, one disjunct each; (1) and (2) are equivalent by Proposition 5.3. For (1) implies (3), suppose t realises the bound and write $A = s_{(n)}(t)$, $B = s_{(n-1)}(t)$, $C = M(n-1)$, $D = \mu_{\mathcal{O}}(n-1)$. We know $A \leq C$ (each of the champion's $n-1$ matches awards at most M), $D \leq B$ (Lemma 5.2) and $A - B = C - D$; over the natural numbers those three force $A = C$ *and* $B = D$ exactly. No positivity is used, so $M = 0$ is not a special case.

Now delete the champion. The remaining $(n-1)$-player tournament has winner's score at least $\mu_{\mathcal{O}}(n-1)$ by definition of the minimum, and at most $B = \mu_{\mathcal{O}}(n-1)$ because deleting a player only lowers scores; so it is exactly $\mu_{\mathcal{O}}(n-1)$, and the player achieving it scored the same before the champion left as after — the champion awarded him nothing. Meanwhile the champion scored $M(n-1)$ from $n-1$ matches each worth at most M, so it took M from every opponent, that one included. Their match is therefore $(M, 0)$ read from the champion's side, and Lemma 5.5(1) places that pair, or its swap, in $\mathcal{O}$. The two formal ingredients are Lemma 5.5(2), which turns "the champion scored $M(n-1)$" into "every one of its matches awarded it M" — a list of $n-1$ entries each at most M and summing to $M(n-1)$ has every entry equal to M — and Lemma 5.5(3), which turns "the best remaining player's score did not drop" into "the second coordinate of that match is 0". No computation enters, and the argument is uniform in n and in $\mathcal{O}$. □

Remark 5.8 (why a disjunction, and what is still open). A result is an ordered pair, and $(0, M)$ is the same match with the two players entered in the other order. One would like to read the condition as "$\mathcal{O}$ contains a total-domination result up to orientation" and be done; that reading is correct, but it is not available as a *simplification* of Theorem 5.7 until one knows that a model and its coordinate-wise reverse induce the same distribution of score sequences, i.e. that $\mathrm{dist}_{\mathcal{O}}(n,\cdot) = \mathrm{dist}_{\mathcal{O}^{\mathrm{rev}}}(n,\cdot)$ for $\mathcal{O}^{\mathrm{rev}} = \{(b,a) : (a,b) \in \mathcal{O}\}$. We have not proved that and do not use it: the two disjuncts are proved separately, and neither is redundant — Proposition 5.15 gives an infinite family attaining the bound with $(M, 0) \notin \mathcal{O}$, and its mirror one attaining it with $(0, M) \notin \mathcal{O}$. The closest published work we found on the reversal question is Thörnblad [18], which extends to generalised tournaments Eplett's criterion for a sequence to be realisable as the score sequence of a tournament isomorphic to its converse; it concerns one tournament against its reverse rather than a distribution over a model.

5.3. **Evaluating the minimax score.** Theorem 5.6 is an exact *reduction*, not a closed form: it turns a question about margins on n players into one about winners' scores on $n-1$ players that does not mention margins. It becomes a closed form exactly where $\mu_{\mathcal{O}}$ is known — which is every model we are aware of in the literature.

Proposition 5.9 (the averaging bound and its tightness)**.** *Let $\mathcal{O}$ be a non-empty outcome model and m a number with $m \le a + b$ for every $(a,b) \in \mathcal{O}$.*

(1) *For every $k \ge 1$, $m(k-1) \le 2\mu_{\mathcal{O}}(k)$: the winner of a k-player tournament scores at least the average. Consequently $M(n-1) - \mu_{\mathcal{O}}(n-1) \le B(M, m, n)$ for every $n \ge 2$ — the sharp bound is never worse than the closed form.*

(2) *If in addition $(a,b) \in \mathcal{O}$ and $(b,a) \in \mathcal{O}$ for some pair with $a + b = m$ and $b \le a \le b+1$, then the inequality in (1) is an equality: $\mu_{\mathcal{O}}(k) = \lceil m(k-1)/2 \rceil$ for every $k \ge 1$, and then $M(n-1) - \mu_{\mathcal{O}}(n-1) = B(M, m, n)$.*

(3) *$\mu_{D_k}(n) = \lceil k(n-1)/2 \rceil$ for every k and every $n \ge 1$.*

Proof sketch. (1) is the total-score bound $m\binom{k}{2} \le \sum_i s_i(t)$ together with $\sum_i s_i(t) \le k\, s_{(k)}(t)$, cancelling $k \ge 1$; it needs no hypothesis beyond m. For the upper bound in (2) one exhibits a tournament whose winner scores $\lceil m(k-1)/2 \rceil$: the near-regular configuration of Proposition 5.12, whose score list is available in closed form. Only $(a,b), (b,a) \in \mathcal{O}$ is needed in the odd case and additionally $b \le a$ in the even one; $a \le b+1$ enters only in the final step identifying the value with the ceiling. (3) is (2) at $m = k$. No computation anywhere. □

Corollary 5.10. *Let $\mathcal{O}$ be an outcome model such that*

(i) *$(M, 0) \in \mathcal{O}$, where $M = M(\mathcal{O})$; and*

(ii) *for some (a,b) with $a + b = m(\mathcal{O})$ and $b \le a \le b+1$, both (a,b) and (b,a) lie in $\mathcal{O}$.*

Then for every $n \ge 2$ the greatest winner's margin realised in $T_n(\mathcal{O})$ is $B(M, m, n)$, with $m = m(\mathcal{O})$.

Proof sketch. Theorem 5.6 under (i), rewritten by Proposition 5.9(2) under (ii). The extrema $M(\mathcal{O})$ and $m(\mathcal{O})$ satisfy the bound hypotheses by definition, and conversely any pair of constants valid for Proposition 5.3 that also satisfies (i) and (ii) must be the extrema, so stating the corollary for the extrema loses nothing. □

Corollary 5.11. *For every k and every $n \ge 2$, the maximum winner's margin in the model D_k is $\lfloor kn/2 \rfloor$; in particular it is $\lfloor n/2 \rfloor$ in the classical model and n in the chess model, at every $n \ge 2$. In the soccer model it is $2n-1$ at every $n \ge 2$.*

Proof sketch. In D_k every result splits exactly k points, so $M = m = k$ and the bound collapses: $k(n-1) - \lceil k(n-2)/2 \rceil = \lfloor kn/2 \rfloor$ for $n \geq 2$, an identity of integers. Hypothesis (i) holds because $(k, 0) \in D_k$, and (ii) because $(\lceil k/2 \rceil, \lfloor k/2 \rfloor)$ and its reverse are both results. For the soccer model, $M = 3$ and $m = 2$, the result $(3, 0)$ gives (i) and the draw $(1, 1)$, being its own reverse, gives (ii), and $3(n-1) - \lceil 2(n-2)/2 \rceil = 2n - 1$. The D_k statement also follows directly from Theorem 5.6 and Proposition 5.9(3), with no mention of hypothesis (ii); in the soccer model the value $\mu_{\mathcal{O}_{\mathrm{soc}}}(k) = k - 1$ that Theorem 5.6 needs comes from Proposition 5.9(2), which is where the draw is used. □

The D_k statement is uniform in k as well as in n. The soccer model is the only one considered here with $M > m$, and its bound is not the bound of any D_k: at $n = 5$ it is 9 while $\lfloor 5k/2 \rfloor$ runs $0, 2, 5, 7, 10, \ldots$ and skips 9.

Proposition 5.12 (the balanced construction)**.** *Assume (ii) of Corollary 5.10 and let (a, b) be the pair it supplies. For every $p \geq 0$ there are tournaments on $2p+1$ and on $2p+2$ players, all of whose matches are settled by (a, b) or (b, a), with score lists*

$$\underbrace{pm,\ \ldots,\ pm}_{2p+1} \qquad \text{and} \qquad \underbrace{pm+a,\ \ldots,\ pm+a}_{p+1},\ \underbrace{pm+b,\ \ldots,\ pm+b}_{p+1}$$

respectively, where $m = a + b$. Adding a player who takes $(M, 0)$ from each of them gives, at every $n \geq 2$, a tournament of margin exactly $B(M, m, n)$ whenever (i) holds as well.

Proof sketch. No computation, and no case analysis on the parity of m: the construction is parameterised by the pair (a, b) rather than by m. The two closed-form score lists come from a single induction on p, through one lemma computing the effect of adding a newcomer whose row is two blocks of equal results to a score list that is itself two blocks of equal values; every row and every intermediate score list has that shape. The largest of these scores is pm in the odd case and $pm + a$ in the even one, and in both parities that is $\lceil m(k-1)/2 \rceil$ on k players. The construction was executed on data before it was formalised: eight models, $n \leq 50$ and $p \leq 24$, against both closed forms, membership in $T_n(\mathcal{O})$ and the closed form of the bound, with no discrepancy. □

We record the upper bound in the closed form too: it was the route taken before the minimax score isolated the content, it is independent of Proposition 5.3, and each checks the other.

Proposition 5.13 (the closed-form upper bound)**.** *Let $\mathcal{O}$ be any outcome model, and let M and m be any numbers such that $a \leq M$ and $b \leq M$ for every $(a, b) \in \mathcal{O}$, and $m \leq a + b$ for every $(a, b) \in \mathcal{O}$. Then for every $n \geq 2$ and every $t \in T_n(\mathcal{O})$ we have $s_{(n)}(t) - s_{(n-1)}(t) \leq B(M, m, n)$.*

Proof sketch. Everything rests on the scalar inequality

$$(\star) \qquad 2\, s_{(n)}(t) + m(n-1)(n-2) \ \leq\ 2 \sum_{i=1}^{n} s_i(t),$$

which says that once a champion is set aside the other $n - 1$ players already carry at least $m\binom{n-1}{2}$ points between them. It is the only thing the argument needs from the direction of Landau's inequality [8], and unlike that inequality it constrains only the maximum against the total, not arbitrary submultisets of the score multiset. It is proved by induction along the seat-one-player recursion, in two cases according to whether the newcomer or a seated player is the champion of the enlarged tournament; in the second case the enlarged maximum is $s + c$, where s is one of the old scores and c is what one entry of the newcomer's row hands to the seated side, and the argument never needs those two to come from the same player, which is where index bookkeeping disappears. From $(\star)$, deleting one occurrence of the maximum and using that a list has sum at most its length times its maximum gives $s_{(n-1)}(t) \geq \lceil m(n-2)/2 \rceil$; with $s_{(n)}(t) \leq M(n-1)$ this is the statement. No computation enters, and the range of n is unrestricted. □

5.4. **Where the bound is and is not attained.**

Proposition 5.14 (the no-draw model). *Let $g \geq 1$ and $\mathcal{O} = \{(g,0),(0,g)\}$: a win is worth g points and a draw is impossible. Then $\mu_{\mathcal{O}}(k) = g\lfloor k/2 \rfloor$ for every $k \geq 1$, and the greatest winner's margin realised on $n \geq 2$ players is $g\lfloor n/2 \rfloor$. For $g \geq 2$ hypothesis (ii) of Corollary 5.10 fails, and the closed form is strictly larger than the truth at every odd $n \geq 3$:*

$$g\left\lfloor \frac{n}{2} \right\rfloor \;<\; B(g,g,n).$$

Proof sketch. The margin value is Theorem 5.6, whose hypotheses hold with $M = g$. For $\mu_{\mathcal{O}}$, the upper bound is Proposition 5.12 at $(a,b) = (g,0)$ — legal because that construction never needs $a \leq b+1$ — and the lower bound needs a parity argument that the averaging bound of Proposition 5.9(1) cannot supply: every result awards multiples of g, hence so does every score, by one induction, and writing $\mu_{\mathcal{O}}(k) = gc$ the averaging bound $g(k-1) \leq 2gc$ gives $c \geq \lfloor k/2 \rfloor$. The strict inequality is then arithmetic. This is an infinite family of models on which the closed form B is provably not the answer, which is what makes dropping hypothesis (ii) a gain rather than a restatement. At $g = 4$, $n = 5$ the truth is 8 and $B(4,4,5) = 10$; over all 2^{10} five-player tournaments the margin marginal is 424 at $d = 0$, 440 at $d = 4$, 160 at $d = 8$ and zero elsewhere — Table 1's classical row stretched by 4, as it must be — and the two bound sequences for $n = 2,\ldots,9$ are $4,4,8,8,12,12,16,16$ against $4,6,8,10,12,14,16,18$. ☐

Proposition 5.15 (the reversed model). *Let $g \geq 1$ and $\mathcal{O} = \{(0,g)\}$, the single result awarding g points to the earlier-seated player of every match. Then $T_k(\mathcal{O})$ has exactly one element, its score list is $0, g, 2g, \ldots, g(k-1)$, so $\mu_{\mathcal{O}}(k) = g(k-1)$, and the greatest realised margin on $n \geq 2$ players is g. Here $(g,0) \notin \mathcal{O}$: the bound is attained through the second disjunct of Theorem 5.7 alone. The mirror model $\{(g,0)\}$ attains it through the first disjunct alone.*

Proof sketch. $|\mathcal{O}| = 1$, so the whole sample space is one tournament and the minimum defining $\mu_{\mathcal{O}}$ is over a single element; its score list is the append-one-match construction of Proposition 5.4(2) iterated. The margin is then $g(n-1) - g(n-2) = g$. As a control, at $n = 5$ the margin of the single tournament of $\{(0,1)\}$ is 1 and neither 0 nor 2 occurs, while the classical model, which *contains* $\{(0,1)\}$, has largest margin 2 there. ☐

Proposition 5.16 (models that never attain the bound). *Both of the following hold at every $n \geq 2$, and both are consequences of Theorem 5.7 rather than of any enumeration.*

(1) *In $\mathcal{O} = \{(1,0),(0,1),(5,5)\}$, where $M = 5$, the bound $5(n-1) - \mu_{\mathcal{O}}(n-1)$ is* not *the greatest realised margin, because neither $(5,0)$ nor $(0,5)$ is a result. At $n = 3$ the bound is 9 and the greatest realised margin is 4.*
(2) *In the all-draws model $\mathcal{O} = \{(1,1)\}$, where $M = 1$, every tournament has margin 0 while the bound $(n-1) - \mu_{\mathcal{O}}(n-1)$ is positive; again neither $(1,0)$ nor $(0,1)$ is a result.*

Neither remaining hypothesis of Theorem 5.7 can be dropped. At $n = 1$ the bound is 0 and the all-draws model attains it while containing neither $(1,0)$ nor $(0,1)$, so $n \geq 2$ is load-bearing; and taking $M = 0$ in the all-draws model — where 0 is not a bound for the results — again makes the bound 0 and attainable while $(0,0)$ is not a result, so the hypothesis that M bounds the model is load-bearing too.

Proof sketch. Each item is Theorem 5.7 instantiated at the model, together with the finite check that the two named pairs are not in a three- or one-element list; the two non-removability statements are single evaluations at $n = 1$ and $n = 2$. The point is the quantifier: before Theorem 5.7 the first item could only be asserted one n at a time by enumerating $T_n(\mathcal{O})$, and it had been recorded only at $n = 3$. ☐

The values read off the computed tables are kept as an independent check on the general theory: they come from the margin marginals of Theorem 4.1, whereas everything above is proved from the definition of a tournament with no table anywhere in its proof.

Theorem 5.17. *The largest winner's margin carrying positive probability on n players is*

$$\left\lfloor \frac{n}{2} \right\rfloor \ \ (\textit{classical},\ 2 \le n \le 13), \qquad n \ \ (\textit{chess},\ 2 \le n \le 9), \qquad 2n-1 \ \ (\textit{soccer},\ 2 \le n \le 7).$$

All three values agree with $B(M, m, n)$ at $(M, m) = (1, 1)$, $(2, 2)$ and $(3, 2)$ respectively.

Proof sketch. For each model and each n in the stated range the margin marginal of Theorem 4.1 is a list of w integers, w being the number of attainable scores, and the assertion is that its largest non-zero index is the stated value: a finite arithmetic check on data already computed, 26 instances in all, re-running no dynamic program. The list is long enough to be conclusive, since a margin is at most the maximum score and so at most $w-1$. The theorem asserts nothing outside the ranges listed. □

Proposition 5.18 (controls on the margin theorems)**.** *Every item below is a count over a complete sample space, computed from the definition of a tournament.*

(1) The bound is off by one in neither direction. *Classically on five players the margin marginal is* $424, 440, 160, 0, 0$*: the bound* $\lfloor 5/2 \rfloor = 2$ *is realised* 160 *times and a margin of* 3 *is impossible. In the chess model on four players it is* $189, 288, 180, 60, 12, 0$*, with bound* 4 *realised* 12 *times. In the soccer model on three players it is* $6, 6, 6, 6, 0, 3, 0$*, with bound* 5 *realised* 3 *times and a hole at* $d = 4$*.*

(2) Hypothesis (i) of Corollary 5.10 cannot be dropped. *In* $\mathcal{O} = \{(1, 0), (0, 1), (5, 5)\}$ *we have* $M = 5$ *and* $m = 1$*; hypothesis (ii) holds, witnessed by the pair* $(1, 0), (0, 1)$*, while neither* $(5, 0)$ *nor* $(0, 5)$ *is a result. On three players* $B(5, 1, 3) = 9$*, while the margin marginal is* $9, 12, 0, 0, 6, 0, \ldots$ *and the greatest realised margin is* 4*.*

(3) Hypothesis (ii) cannot be dropped either. *In* $\mathcal{O} = \{(4, 0), (0, 4)\}$ *we have* $M = m = 4$ *and (i) holds, while (ii) fails: the only pair* (a, b) *with* $a + b = 4$ *and* $b \le a \le b + 1$ *is the draw* $(2, 2)$*, which is not a result — checked by quantifying over all candidate pairs, not asserted for one. On three players* $B(4, 4, 3) = 6$*, while the margin marginal is* $2, 0, 0, 0, 6, 0, 0, 0$ *and the greatest realised margin is* 4*.*

(4) Hypothesis (ii) is not necessary for attainment. *In* $\mathcal{O} = \{(3, 0), (0, 3), (2, 1)\}$ *we have* $M = m = 3$ *and (i) holds, and (ii) fails because* $(1, 2)$ *is not a result; yet* $B(3, 3, 3) = 4$ *and the margin marginal on three players is* $4, 8, 4, 8, 3, 0$*, so the bound is attained.*

(5) The closed form reproduces the computed values. $B(1, 1, n) = \lfloor n/2 \rfloor$ *for* $2 \le n \le 13$*,* $B(2, 2, n) = n$ *for* $2 \le n \le 9$ *and* $B(3, 2, n) = 2n - 1$ *for* $2 \le n \le 7$ *— the right-hand sides of Theorem 5.17, which were obtained from the tables and not from any formula.*

Proof sketch. Items (1)–(4) are exhaustive enumerations of complete sample spaces, each built from the definition of a tournament with the dynamic program nowhere in the evaluated expression: 2^{10} classical tournaments on five players, 3^6 chess tournaments on four players, and 27, 27, 8 and 27 tournaments on three players in the soccer model and in the three models named in (2)–(4). Item (5) is finite arithmetic on the closed form over the 26 values of n listed. Note that in all three of the models in (2)–(4) the upper bounds of Propositions 5.3 and 5.13 still hold, as they must, since neither assumes either hypothesis. □

Remark 5.19 (the two hypotheses, read through the minimax score). This section explains the three models of Proposition 5.18(2)–(4), which were previously recorded as bare enumerations. In (2) hypothesis (i) fails in the strong sense of Theorem 5.7 — neither orientation of the total-domination result is present — so the sharp bound, which is also 9 there, is missed at *every* n, not only at $n = 3$. In (3) hypothesis (i) holds, so the greatest realised margin is the sharp bound $4 \cdot 2 - \mu_{\mathcal{O}}(2) = 4$; what is missed is the closed form $B(4, 4, 3) = 6$, and Proposition 5.14 identifies that model as the case $g = 4$ of an infinite family where the closed form overshoots at odd n. In (4) (i) holds too, so the sharp bound is attained at every n; the *closed form* is attained as well at $n = 3$ and $n = 4$, because there the minimax score meets the averaging bound of Proposition 5.9(1) anyway at $k = 2, 3$, so the two bounds coincide. Hypothesis (ii) was never a condition on attainment: it is a sufficient condition for $\mu_{\mathcal{O}}(k) = \lceil m(k-1)/2 \rceil$.

Provenance of the statements. A novelty claim is worth only the search behind it, so we record ours; it was run in four passes over the same three instruments — Moon's monograph [16], a local full-text corpus of arXiv mathematics sources (about 85 gigabytes in 367 shards, all files scanned), and the arXiv API and [17] live, always with a positive control in the same run.

Pass 1. Moon's §24 is titled *The Largest Score in a Tournament*; we read it in full, and the phrases "second largest score", "second highest score", "second maximum score" and even "maximum score" do not occur anywhere in the book (2013 electronic edition, full text searched). In the corpus, "second maximum score" returned nothing and "second maximal score" exactly one occurrence, in [13] itself; further passes for "winning margin" with tournament context, for the gap between the two largest scores, and for $\lfloor n/2 \rfloor$ with tournament context returned nothing relevant. Seventeen live arXiv queries, control "round-robin tournament" returning 58 hits, returned nothing for the margin, and OEIS name searches for the margin, the runner-up score and co-champions returned nothing.

Pass 2, with Corollary 5.10: fifteen phrases in eight groups over the corpus — among them "gap between the two largest scores" (0), "winner's margin" (2, neither about tournaments), "largest winning margin", "extremal score sequence" and "maximum gap between scores" (0 each) — ten live arXiv queries, one of them a control that returned all 8 expected hits, and six OEIS searches with two controls that both returned. One query, "maximum score" with "round-robin tournament", returned [15] and [1], both about the *uniqueness* of the maximum score. We read the latter in full: "margin" occurs in it only inside page-layout commands, and "second largest", "second maximum", "gap" and "runner" not at all; it is prior art for Remark 2.4 and for nothing in this section.

Pass 3, with Theorem 5.6, targeting the minimax score: twenty-one corpus phrases, thirteen live arXiv queries with a control returning all 20 hits, five OEIS searches with two controls. "Minimax score" occurs 37 times and is a term collision in both directions — every occurrence is the Minimax *voting rule* of computational social choice or "minimax score risk" in statistics, none a tournament score sequence. "Minimum maximum score", "smallest possible maximum score", "minimum over all tournaments", "minimum winning score" and "least maximum score" returned 0 each; "minimax score" with "tournament" returned 0 on arXiv, and "countable-outcome" returned four papers, two of them [13] and [1] and two unrelated. "Near-regular tournament" does return 21 corpus hits, and that literature is exactly the classical case $\mu_{D_1}(k) = \lceil (k-1)/2 \rceil$ of Proposition 5.9(3), which we do not claim.

Pass 4, with Theorem 5.7: twenty-one corpus phrases, eleven live arXiv queries with a control returning all 20 hits, three OEIS searches with two controls. "Maximum score difference" occurs 17 times, all voting or machine learning; "largest" and "maximal score difference" 4 each, likewise; "margin is attained" 5 times, all support-vector margins; "attains the maximum margin" and "maximum margin is achieved" 0; "dominant player" 142 times, always in the game-theoretic sense rather than a player who beats everyone, and "dominant vertex" 442 times in the sense of the domination number. The arXiv query for "score sequence", "necessary and sufficient" and "tournament", all three in the abstract, returned five papers, whose abstracts we read in full: Iványi [6], whose criterion for a multigraph to admit an orientation with a prescribed out-degree vector has the theorems of Landau and Moon as special cases; Kolesnik and Sanchez [7], short proofs of those same two theorems by zonotopes; Thörnblad [18], Eplett's criterion for the score sequences of self-converse tournaments, extended to generalised tournaments; Liu [9], on reconstructing a score sequence from a score set; and Guo, Juillet and Tang [5], on explicit constructions of the football model whose existence gives a short proof of Moon's theorem. All five belong to the Landau–Moon *realisability* question — which sequences are score sequences — and none of them mentions the winner's margin, let alone when an extremal one is realised. Nothing found in any pass states either the sharp bound or the characterisation. We note that the question Theorem 5.7 answers is one the countable-outcome framing makes askable at all: in the classical model $(1, 0)$ is always a result, so the hypothesis is invisible there.

Remark 5.20 (a term collision). "Margin of victory" is a taken term in computational social choice, where it denotes the smallest number of votes that must be changed to change the outcome of an election [19], and where it has been carried over to tournament solutions — the winner sets of voting rules on the majority tournament — by a line of papers including one titled *Margin of Victory in Tournaments* [2]. That work concerns voting rules on tournaments, not score sequences, and is not prior art for anything here; but the collision is total — every one of the 775 occurrences of "margin of victory" in the corpus we searched is of that kind — so the phrase should be disambiguated wherever this result is quoted, and we have used "winner's margin" throughout.

6. Controls

A table of large integers is worth exactly as much as the checks that would have caught it being wrong. Seven groups were run; all are part of the formal development, and the last of them — an external check against a published sequence — gets its own subsection.

Proposition 6.1. (1) The definitions mean what their names say. *On* $(3,3,1)$ *the runner-up score is* 3*, the margin* 0 *and the number of co-champions* 2*; on* $(2,1,0)$ *they are* 1*,* 1*,* 1*; on* $(2,2,2)$ *they are* 2*,* 0*,* 3*.*

(2) The length hypothesis of Theorem 3.3 is not removable*: on* (0) *the maximum is attained once while the margin is* 0*; on* $(1,1)$ *and* $(2,0)$ *the equivalence is non-vacuous in both directions.*

(3) The joint matrix determines both marginals. *Its row sums are the winner's-score marginal and its column sums the runner-up marginal, checked against the computed table at classical* $n=4,\dots,10$ *and chess* $n=4,\dots,7$*.*

(4) Everything sums to the sample space. *Each of the three marginals, and the joint matrix, sums to* $|\mathcal{O}|^{\binom{n}{2}}$ *at classical* $n \le 13$ *and chess* $n \le 9$*; and at* $n=9$ *and* $n=12$ *the margin-zero entry plus the* $k=1$ *co-champion entry is* 2^{36} *resp.* 2^{66}*, which is Corollary 3.4 read numerically off two different components of the same computation.*

(5) Too-large claims are refuted. *Two players cannot both win every match, so the joint cell* $(a,b)=(5,5)$ *on six players is* 0 *while the cell* $(5,4)$ *is* 1920 *and is not* 1921*; a margin of* 7 *on twelve players is impossible, while a margin of* 6 *occurs* $579\,018\,101\,760$ *times.*

(6) The sample space is enumerated directly. *For classical* $n=5,6$ *and chess* $n=4$ *all four statistics were recomputed by building every one of the* 2^{10}*,* 2^{15} *and* 3^6 *tournaments from the definition, computing each score list and tallying, with the dynamic program nowhere in the evaluated expression. The results agree with the table. This is the control that would catch a wrong score function, a wrong enumeration or a wrong bridge, so the numbers do not rest on Proposition 2.2 alone.*

(7) The co-champion distribution is not monotone in k. *On seven players no tournament has exactly six co-champions, while* $2\,640$ *have seven — all* n *players tie exactly when the tournament is regular, which forces* n *odd, and* $2\,640$ *is the number of labelled regular tournaments on seven vertices (A007079). An interior zero is a fact about the distribution and it refutes the natural guess that the counts decrease in* k*.*

Proof sketch. Items (1), (2), (4) and (5) are finite arithmetic on data already stated and are checked by the kernel with no compiled evaluation. Items (3), (6) and (7) re-run finite computations: for (3), 11 instances comparing two length-w vectors; for (6), the exhaustive enumerations of 1024, 32 768 and 729 tournaments named there; for (7), four evaluations of the co-champion statistic on the tables at $n=5,6,7$. □

Two further groups belong to Section 5. The minimax score was computed by direct enumeration in six instances — $\mu_{D_1}(5)=2$ over 1024 tournaments, $\mu_{D_2}(4)=3$ over 729, $\mu_{\mathcal{O}_{\mathrm{soc}}}(3)=2$ over 27, and three more — each compared against the closed form Proposition 5.9 predicts. And Lemma 5.5 was pinned down on the cyclic three-player classical tournament, whose players' match lists are $(1,0),(0,1)$; $(0,1),(1,0)$; and $(1,0),(0,1)$: the earliest-seated player reads both of its matches off the other players' rows and so sees both of them reversed, while the latest-seated player reads both off its own row. Classically the reversed pairs are results as well, so it takes a one-sided model to see that the disjunction in Lemma 5.5(1) cannot be dropped: in $\mathcal{O}=\{(0,1)\}$ the earliest-seated of three players has both matches recorded as $(1,0)$, which is not a result, while its swap is.

6.1. **The one published constraint beyond the first column.** The co-champion distribution $T(n,k)$ of Table 2 has exactly two columns that the literature pins. Its first, $T(n,1)$, is the unique-winner count A013976, known to $n = 20$; its last, $T(n,n)$, counts regular tournaments and is A007079 for odd n and 0 for even n. For the interior there is one published quantity: sequence A125031 of [17] — contributed by M. Fuller in 2006, with terms to $n = 16$ in the entry and a *b*-file to $n = 20$ due to A. Howroyd — the total number of highest scorers over all $2^{n(n-1)/2}$ tournaments on n players, which is the first moment $\sum_k k\,T(n,k)$.

Proposition 6.2. *For every n from 2 to 13, $\sum_k k\,T(n,k)$ equals A125031(n):*

$$\begin{array}{rlrl} n=2: & 2 & n=8: & 430\,790\,144 \\ n=3: & 12 & n=9: & 111\,823\,251\,840 \\ n=4: & 104 & n=10: & 56\,741\,417\,927\,680 \\ n=5: & 1\,560 & n=11: & 57\,729\,973\,360\,342\,272 \\ n=6: & 53\,184 & n=12: & 118\,195\,918\,779\,085\,344\,768 \\ n=7: & 3\,422\,384 & n=13: & 479\,770\,203\,506\,298\,422\,135\,808. \end{array}$$

Moreover the entry $T(4,k)$ for $k = 1,2,3$ is $32, 24, 8$, which is the row published in that sequence's own example field.

Proof sketch. A finite arithmetic identity on the twelve co-champion vectors of Theorem 4.1, checked by the kernel; the right-hand sides are the published terms. □

This is a genuine external check on the part of the co-champion distribution that no other published quantity touches, and it is not implied by the score-sequence tables being right: a compensating error across two columns of $T(n,\cdot)$ would have to preserve the first moment at all twelve values of n.

7. The soccer model

The correctness proof of Proposition 2.2 is stated for an arbitrary outcome model, and the constant match total of the family D_k is used nowhere in it. A model in which the two players' scores do *not* sum to a constant therefore costs nothing.

Proposition 7.1. *Let $\mathcal{O}_{\text{soc}} = \{(3,0),(1,1),(0,3)\}$. The two players' scores in one match sum to 3, 2, 3 according to the result, so $\mathcal{O}_{\text{soc}} \neq D_k$ for every k and the model lies outside the setting $X_{ij} + X_{ji} = 1$ of* [13]. *Every statement of Sections 2–3 applies to it unchanged, and for $2 \le n \le 7$ the winner's-score marginal, the number of realisable score vectors, the unique-winner count and the four top-two statistics are all known exactly. In particular the numbers of realisable score vectors are $2, 7, 40, 355, 3\,678, 37\,263$ and the unique-winner counts are*

$$2, \quad 21, \quad 576, \quad 46\,785, \quad 11\,618\,964, \quad 8\,559\,181\,806$$

for $n = 2,\dots,7$, each row of the underlying table summing to $3^{\binom{n}{2}}$.

Proof sketch. $\mathcal{O}_{\text{soc}} \neq D_k$ is immediate: matching lengths forces $k = 2$, and the two lists then differ. The values are one evaluation of the dynamic program per n, exactly as in Theorem 4.1, with $w = 3n-2$ attainable scores; the largest instance, $n = 7$, has $37\,263$ states. □

The soccer instance also serves as the strongest single external check in this paper, because it is a model the development had never touched and a published table produced by an independent implementation. Zeilberger's soccer output [3] prints the full table for $n = 3,\dots,7$. We parsed all five tables and compared them cell by cell against ours: $41\,343$ **score vectors, no mismatch**, with per-n totals exactly $3^{\binom{n}{2}}$.

Soccer $n = 8$ was not attempted; a count outside the proof assistant gives $361\,058$ realisable score vectors there, which extrapolates past the memory we allow one computation, and [3] stops at $n = 7$ as well.

8. Reproduction of the published material

For completeness we record what the same machinery reproduces of the classical literature, since the tables of Section 4 are read off it.

Proposition 8.1. *The full exact distribution of the score sequence is reproduced for the classical model at $n \leq 10$ and the chess model at $n \leq 7$ — in particular over MacMahon's two ranges, classical $n \leq 9$* [10] *and chess $n \leq 6$* [11]*, and one step past each. Every cell agrees with the tables of* [3] *(classical $n \leq 12$, 14 805 cells at the top; chess $n \leq 9$, 26 376 cells).*

Proposition 8.2. *The winner's-score marginal, the unique-winner count and the number of realisable score sequences are known exactly for classical $n \leq 13$ and chess $n \leq 9$. The classical unique-winner counts are A013976 and agree at every n; the chess ones are 2, 21, 540, 42 765, 10 427 004, 7 639 657 074, 16 774 320 322 584, 110 539 053 095 823 189 for $n = 2, \ldots, 9$. The classical case $n = 14$ was computed outside the proof assistant and is not asserted here.*

Proposition 8.3. *The frequency of the nine-player classical score sequence $(2, 2, 3, 3, 4, 4, 6, 6, 6)$, and of its complement $(2, 2, 2, 4, 4, 5, 5, 6, 6)$, is 361 307 520 and is not the value 361 297 520 printed in* [10].

The erratum is due to Malinovsky and Moon [15, §1], who found it by observing that the frequency must be divisible by 9 — in either sequence there are nine choices of label for the winner of the match between the two players of score 4 — and that the printed value is not; the two frequencies each need to be increased by 10 000. Our computation, made from the definition over the whole 2^{36}-element sample space, gives 361 307 520 for both vectors, an independent confirmation of their correction.

Proposition 8.4. *Landau's criterion* [8] *characterises the support: over every non-decreasing vector of length n with entries at most $n-1$ — 1716 of them at $n = 7$ — "the computed count is non-zero" and "Landau's condition holds" agree, for $n = 5, 6, 7$, and the number of admissible vectors is 9, 22, 59, i.e. A000571. Further controls: no vector failing Landau's condition or having the wrong total receives a positive count; every tabulated cell is strictly positive, so the development does not prove everything to be zero; the state cannot be coarsened to the score total, since $(0, 2)$ and $(1, 1)$ have equal sums and different one-player extensions; the two players' scores in a match of D_k sum to k; and the enumeration $T_n(\mathcal{O})$ really does have $|\mathcal{O}|^{\binom{n}{2}}$ elements.*

9. Verification

Every mathematical claim in this paper is formally verified in Lean 4 against *Mathlib* — the two counts flagged in Sections 7 and 8 as computed outside the proof assistant excepted; the development contains no `sorry`, and the repository reference is to be added at submission. The reasoning layer is kernel-clean, depending on no axioms beyond `propext`, `Classical.choice` and `Quot.sound`. It covers Proposition 2.2 and the transfer principle Proposition 2.3, the permutation invariance of the four statistics and the inequality $s_{(n-1)} \leq s_{(n)}$ (Proposition 3.2), Theorem 3.3, Corollary 3.4, the identification of the co-champion distribution's first column with the unique-winner count (which holds by definitional unfolding), and the statement that the soccer model is not of the form D_k. It also covers the whole of Section 5's general theory: every lemma, proposition, theorem and corollary there, from Lemma 5.2 to Proposition 5.16, is proved from the definition of a tournament with no table and no compiled evaluation in its proof. Where a statement quantifies over models — "every D_k", "every no-draw model", "every reversed model" — the quantifier is in the formal statement, not in a family of instances.

What is delegated to compiled evaluation is exactly the finite computations of the tables: one evaluation of the dynamic program per model and per n for Theorem 4.1 and Proposition 7.1, and the direct enumerations of the sample space in Proposition 6.1(6). Each theorem stating such a computed value carries, beside `propext` and `Quot.sound`, exactly one further axiom: the one Lean's native-evaluation bridge generates to record that the compiled code returned the stated value. The restatements about $T_n(\mathcal{O})$ add `Classical.choice` from the kernel-clean bridge and nothing else.

Theorem 5.17, Proposition 6.2 and items (1), (2), (4), (5) of Proposition 6.1 are decided by the kernel on the already-stated data, so their only computational dependence is on the tables they read; Propositions 5.18 and 5.16 and the enumerated minimax values are decided by the kernel as well, but on complete sample spaces rebuilt from the definition, so they depend on no table and on no theorem of Section 5. No result of Section 5 uses compiled evaluation at all.

The largest single evaluation, classical $n = 13$, took about four minutes and 5.6 gigabytes. The six modules carrying Section 5 elaborate in 11, 15, 23, 50, 22 and 8–16 seconds of wall-clock time — the last measured twice, page-cache dependent — and add between a quarter and one gigabyte each over the cost of loading the library. Outside the proof assistant, every value was first produced by an independent implementation in Python, and every cell of the underlying score-sequence tables was compared against the published tables of [3] over their range, cell by cell, before any formal statement was written — 14 805 classical cells at $n = 12$, 26 376 chess cells at $n = 9$ and 41 343 soccer cells over $n = 3, \ldots, 7$.

Section 5 was gated the same way, in three rounds, each in a reimplementation of the formal definitions written one definition at a time so that a check there is a check of the formal statement rather than of a paraphrase. Round one ran the complete sample space of nine outcome models at every $n \leq 6$ with $|\mathcal{O}|^{\binom{n}{2}} \leq 4 \cdot 10^6$, checking $(\star)$, the two one-line bounds, the runner-up inequality, the upper bound — which held everywhere — and attainment, which held in seven of the nine models; the induction step of $(\star)$ as an abstract statement about lists on 30 000 random inputs, of which 24 384 satisfied its hypotheses, with no failure; the witness of Proposition 5.12 in eight models for $n \leq 50$; and $B(k, k, n) = \lfloor kn/2 \rfloor$ for $k \leq 8$, $n \leq 59$. Round two, for Theorem 5.6, checked the deletion of Lemma 5.2 and the runner-up bound over every (tournament, deleted player) pair in ten models with $|\mathcal{O}|^{\binom{n}{2}} \leq 6 \cdot 10^4$; the value $\max_t \operatorname{gap}(t) = M(n-1) - \mu_{\mathcal{O}}(n-1)$ in ten models at $n = 2, \ldots, 7$ with $|\mathcal{O}|^{\binom{n}{2}} \leq 4 \cdot 10^6$, exact in all nine models containing a total-domination result and missed only by the one that does not; the averaging bound and its tightness for $k = 1, \ldots, 7$; and the six enumerated minimax values. Round three, for Theorem 5.7, checked the three items of Lemma 5.5 over *every* (tournament, player) pair in twelve models with $|\mathcal{O}|^{\binom{n}{2}} \leq 6 \cdot 10^4$; both directions of the characterisation in twelve fixed models at $n = 2, \ldots, 7$ and at every legal value of M; the append-one-match construction of Proposition 5.4(2) in ten models at $n = 2, \ldots, 6$; and both directions again on 4 000 randomly generated models with $|\mathcal{O}| \leq 4$ and entries at most 3, at $n = 2, 3, 4$ and two values of M each — 24 000 triples, of which the bound is attained in 7 470 and missed in 16 530, with no failure either way. That split is what makes the check discriminating rather than vacuous: two thirds of the random triples do not attain, and in every one of those the model contains neither $(M, 0)$ nor $(0, M)$.

References

[1] G. Amir and Y. Malinovsky, *Uniqueness of maximum scores in countable-outcome round-robin tournaments*, Discrete Applied Mathematics **384** (2026) 16–22; arXiv:2411.02141. The result restated in Remark 2.4 is its main result, read from the arXiv version v3 (3 December 2025).

[2] M. Brill, U. Schmidt-Kraepelin and W. Suksompong, *Margin of Victory in Tournaments: Structural and Experimental Results*, Proceedings of the AAAI Conference on Artificial Intelligence **35**(6) (2021) 5228–5235.

[3] S. B. Ekhad (D. Zeilberger), *On the Most Commonly-Occurring Score Vectors of American Tournaments of n-players, and their Corresponding Records*, Personal Journal of Shalosh B. Ekhad and Doron Zeilberger, 13 June 2016 (version of 26 July 2016); `https://sites.math.rutgers.edu/~zeilberg/mamarim/mamarimhtml/Percy.html`, with the output files `oPERCYG0.txt` (classical, $n \leq 12$), `oPERCYG1.txt` (chess, $n \leq 9$) and `oPERCYG2.txt` (soccer, $n \leq 7$).

[4] R. A. Epstein, *The Theory of Gambling and Statistical Logic*, first edition, Academic Press, New York, 1967.

[5] G. Guo, N. Juillet and W. Tang, *The football model, stochastic ordering and martingale transport*, Bernoulli **32**(3) (2026) 2323–2350; arXiv:2503.07145.

[6] A. Iványi, *Reconstruction of complete interval tournaments*, Acta Universitatis Sapientiae, Informatica **1**(1) (2009) 71–88; arXiv:1003.4016.

[7] B. Kolesnik and M. Sanchez, *The geometry of random tournaments*, Discrete & Computational Geometry **71**(4) (2024) 1343–1351; arXiv:2009.09322 (where the two authors are listed in the other order).

[8] H. G. Landau, *On dominance relations and the structure of animal societies. III. The condition for a score structure*, Bulletin of Mathematical Biophysics **15** (1953) 143–148.

[9] B. Liu, *New Theoretical Insights and Algorithmic Solutions for Reconstructing Score Sequences from Tournament Score Sets*, arXiv:2512.16961v1 (2025).

[10] P. A. MacMahon, *An American tournament treated by the calculus of symmetric functions*, Quarterly Journal of Pure and Applied Mathematics **49**, no. 193, 1–36; reprinted in *Percy Alexander MacMahon: Collected Papers*, Vol. I (G. E. Andrews, ed.), MIT Press, 1978, pp. 308–343. (Volume 49 appeared in parts over 1920–1923; this paper is dated 1920 in [3] and [17] and 1923 in [13] and [15].)

[11] P. A. MacMahon, *Chess tournaments and the like treated by the calculus of symmetric functions*, Quarterly Journal of Pure and Applied Mathematics **49**, 353–384; reprinted in *Collected Papers*, Vol. I, pp. 344–375.

[12] Y. Malinovsky, *Distinct Extreme Scores in Random Round-Robin Tournaments*, arXiv:2603.06886 (2026).

[13] Y. Malinovsky, *Extreme Score Distributions in Countable-Outcome Round-Robin Tournaments of Equally Strong Players*, Methodology and Computing in Applied Probability **28**(2) (2026), paper no. 58; arXiv:2601.15950. Quotations and the numbering of Corollary 2 are from the arXiv version v3 (20 May 2026); the two sentences quoted here read identically in v1 (22 January 2026).

[14] Y. Malinovsky, *On the distribution of winners' scores in a round-robin tournament*, arXiv:2201.05018 (2022).

[15] Y. Malinovsky and J. W. Moon, *On round-robin tournaments with a unique maximum score*, Australasian Journal of Combinatorics **89**(1) (2024) 24–31; arXiv:2208.14932.

[16] J. W. Moon, *Topics on Tournaments*, Holt, Rinehart and Winston, New York, 1968; electronic edition, Project Gutenberg eBook 42833, 2013, produced with the author and carrying minor typographical corrections together with a revised proof of its Theorem 34.

[17] OEIS Foundation Inc., *The On-Line Encyclopedia of Integer Sequences*, `https://oeis.org`; entries A000571, A007079, A007747, A013976, A125031, A274098.

[18] E. Thörnblad, *Eplett's theorem for self-converse generalised tournaments*, Australasian Journal of Combinatorics **70**(3) (2018) 329–335; arXiv:1606.02081.

[19] L. Xia, *Computing the margin of victory for various voting rules*, Proceedings of the 13th ACM Conference on Electronic Commerce (EC '12), 2012, pp. 982–999.