Do all Tournaments Admit Irrelevant Matches?

Marco Faella*
University of Naples Federico II
Naples, Italy
m.faella@unina.it

Luigi Sauro
University of Naples Federico II
Naples, Italy
luigi.sauro@unina.it

ABSTRACT

We consider tournaments played by a set of agents in order to establish a ranking among them. We introduce the notion of irrelevant match, as a match that does not influence the ultimate ranking of the involved parties. After discussing the basic properties of this notion, we seek out tournaments that have no irrelevant matches, focusing on the class of tournaments where each agent challenges each other exactly once. We prove that tournaments with a static schedule and at least 5 agents always include irrelevant matches. Conversely, dynamic schedules can be devised in ways that avoid irrelevant matches, at least for one of the involved agents.

KEYWORDS

Tournaments; Game theory for practical applications; Social choice theory

ACM Reference Format:
Marco Faella and Luigi Sauro. 2018. Do all Tournaments Admit Irrelevant Matches?. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018,* IFAAMAS, 8 pages.

1 INTRODUCTION

Tournaments are sets of pairwise contests, called matches, aimed at establishing a ranking among a set of participants. Sport competitions are often organized in tournaments that attract significant popular interest and financial resources. Depending on the type of tournament and on the ranking rule, there may be matches that have no effect on the ultimate ranking of the two contestants. For instance, in 2014 the last match in the preliminary group stage of the UEFA Champions League (Group H) was Porto vs Šachtar. Before the match, Porto was on top of the group with 13 points whereas Šachtar was second with 8 points. Since in football tournaments winning a match yields 3 points, Šachtar had no chance of overtaking Porto, so any possible outcome of the match would not have changed the final ranking. Incidentally, the match ended in an unexciting 1–1 tie.

Clearly, such irrelevant matches are undesirable, as they lack the incentive that is the very essence of a tournament. In this paper, we formalize the notion of relevant and irrelevant match and we tackle the problem of designing tournaments that contain no irrelevant matches.

The very definition of relevance is not trivial and it is the subject of Section 3. Assume for simplicity that a match between a and b can only end in two possible ways: a wins or b wins. We start from the intuition that a match is relevant for one of the involved agents if the future ranking of that agent depends on its outcome. If the future ranking is uniquely determined by the outcome of the present match (say that we are analysing the last match in the tournament), we only need to compare two rankings from the point of view of our agent. We stipulate that agents compare two rankings according to the following preference: they prefer to have fewer agents above them; equal that, they prefer to have fewer agents at the same level as them.

Generally, the future ranking is not uniquely determined by the outcome of the present match. Other matches may follow, possibly influencing the ultimate ranking. We resolve the uncertainty on the future matches via a probabilistic setting, in which agents hold prior beliefs on the relative strengths of their peers. As in the classical Bradley-Terry model [1], such beliefs can be used to estimate the probability that any future match ends in either way. We then define a match to be relevant for an agent if the ranking distribution corresponding to one of the two outcomes *stochastically dominates* the ranking distribution induced by the other outcome [4] (Definition 3.2). We also develop an equivalent non-probabilistic characterization of relevance for the class of point-based round-robin tournaments [1].

Once the basic definitions have been laid out, in Section 4 we study the circumstances under which relevant tournaments exist. In common tournaments, the sequence of matches is fixed *a-priori*, regardless of the outcomes. We call these *static* tournaments. We also consider the more general class of *dynamic* tournaments, in which the sequence of matches adapts according to past outcomes. We then obtain the following main results:

(1) All round-robin tournaments with a static schedule and at least 5 agents include a match that is irrelevant for both of the involved agents (Theorem 4.1).
(2) For all numbers of agents, there is a round-robin tournament with a dynamic schedule where all matches are relevant for at least one of the involved agents (Theorem 4.7).
(3) All round-robin *balanced* tournaments with at least 6 agents include a match that is irrelevant for one of the involved agents (Theorem 4.8).

The present preliminary work could be extended in various directions, some of which are described in our conclusive Section 5.

Related work. In graph theory, a tournament is a complete and asymmetric directed graph [9]. Nodes represent agents and an edge from a to b indicates that a won the match against b. In common

*This author was partially supported by Istituto Nazionale di Alta Matematica (INdAM).

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden.

[1] In a round-robin tournament each agent challenges all other agents exactly once. A tournament is point-based if agents accrue points corresponding to winning or losing each match, and the final ranking is based on such points [10].

terms, a tournament is in fact an *outcome* of a round-robin (real-world) tournament. Hence, graph-theoretic tournaments ignore the temporal aspects and only focus on the ultimate outcome of all the matches. On the other hand, a match may be relevant if played at the beginning of a tournament and irrelevant if played towards the end. So, in this paper we explicitly model the temporal ordering of the matches.

The problem of ranking the participants to a tournament has been studied in the literature, with the two most prominent approaches being the maximum likelihood method and the points system [7, 10]. In this paper, we focus on the points system, as it is commonly used in real-world tournaments.

Tournament rankings are special cases of rank aggregation problems, widely studied in machine learning [2, 13]. The common objectives in that area are orthogonal to ours, mainly concerning accuracy (i.e., obtaining a ranking that is close to some model of ground truth) and efficiency (i.e., obtaining a ranking from few comparisons).

Finally, the Economics literature discusses of tournament design with the objective of maximizing profit for the organizers [3, 12]. The inquiry that is most related to ours finds strong correlation between match attendance and *importance* for either agent, where the latter is measured by an ad-hoc formula based on the possibility that the agent will win the championship and the number of remaining matches [5].

To the best of our knowledge, the notion of (ir)relevant match was not considered in the literature.

2 PRELIMINARIES

Informally, by tournament we mean a schema of possible matches among multiple agents which, according to the outcomes of the single matches, eventually returns a scoreboard. More specifically, in this paper we focus on tournaments that are *ordinal* in the sense that what really matters is the mutual placement of agents rather than their absolute scores. For this reason, we assume that the final outcome of a tournament is a *ranking* among agents which, by allowing possible ties, is represented by a weak order.

In the following, unless differently specified, we assume a fixed set of agents $A = \{1, \ldots, n\}$. A *ranking* is a weak order on A, i.e., a total, reflexive, and transitive relation, $\preceq \subseteq A \times A$. As usual, by $\sim$ and $\prec$ we denote the symmetric and asymmetric parts of $\preceq$, respectively. Intuitively, $a \prec b$ means that the agent b has a better placement than a, whereas $a \sim b$ means that they are ranked the same (a.k.a. a *tie*).

We call *match* a pair of distinct agents. Generally speaking, a match may have multiple outcomes. Here, we focus on binary tournaments, whose matches can have one of two possible *outcomes* w or l.

Then, a *tournament* $\mathcal{T}$ is a labelled full binary tree where: *(i)* each internal node is labelled with a match; *(ii)* the left (resp., right) arc coming out of an internal node is labelled with the outcome w (resp., l); *(iii)* each leaf is labelled with a ranking. Intuitively, given an internal node x labelled with a match (a, b), the left child x_{w} corresponds to a having won the match and the right child x_{l} corresponds to a having lost the match (and therefore, b having won it).

The labels attached to a node and to either one of its outgoing arcs describe a match and its result, respectively. We collect these labels in a triple (a, b, o), called an *event*. A path from an internal node to a leaf induces a sequence of events and a final ranking, denoted by the sequence $\langle e_1, \ldots, e_h, \preceq\rangle$. When this does not cause confusion, we will identify a path and its sequence of events. We call *full path* a path from the root of $\mathcal{T}$. Given a path π, we denote by $\mathsf{won}(\pi, a)$ (resp. $\mathsf{lost}(\pi, a)$) the number of matches won (resp. lost) by a in π.

Types of tournaments. We distinguish the following families of tournaments:

A tournament is *static* if it is a complete tree and all its internal nodes at a given level are labelled with the same match. In other words, the sequence of matches is the same on all full paths.

A *round-robin* tournament is a tournament where in all full paths each agent challenges all the other agents exactly once. Notice that each full path in a round-robin tournament with n agents comprises $\binom{n}{2} = \frac{n(n-1)}{2}$ events.

A tournament is *point-based* if agents implicitly accrue points corresponding to winning or losing each match, and the final ranking is based on such points. Formally, let

$$\mathsf{score}(c, \pi) = \mathsf{won}(c, \pi)p_{\mathsf{w}} + \mathsf{lost}(c, \pi)p_{\mathsf{l}}\,,$$

where c is an agent, π is a full path, and $p_{\mathsf{w}}, p_{\mathsf{l}}$ are real numbers. We say that a tournament is point-based if there exist $p_{\mathsf{l}} < p_{\mathsf{w}}$ such that for all full paths $\pi = \langle e_1, \ldots, e_h, \preceq\rangle$ and agents a and b, we have that $a \preceq b$ iff $\mathsf{score}(a, \pi) \leq \mathsf{score}(b, \pi)$.

A round-robin tournament is *balanced* if n is even and for all full paths $\pi = \langle e_1, \ldots, e_h, \preceq\rangle$ and all $i = 0, \ldots, n-2$, the set of $\frac{n}{2}$ events

$$\left\{ e_j \,\middle|\, j = i\frac{n}{2} + 1,\ i\frac{n}{2} + 2,\ \ldots,\ (i+1)\frac{n}{2} \right\}$$

involves all agents. In other words, balanced tournaments are organized into $n-1$ *rounds*. Every agent plays one match in each round.

Example 2.1. Figure 1 shows a tournament $\mathcal{T}$ among the agents a, b, and c. In particular, the path highlighted in bold corresponds to $\pi = \langle (a, b, \mathsf{w}), (a, c, \mathsf{l}), (b, c, \mathsf{w}), \preceq\rangle$, where in the final ranking $\preceq$ all the agents tie.

Note that all full paths from the root to a leaf share the same sequence of matches where each agent challenges the other ones exactly once. Consequently, $\mathcal{T}$ is a static round-robin tournament. Moreover, the placement of an agent at the end of a full path $\pi = \langle e_1, e_2, e_3, \preceq\rangle$ depends on the number of matches it won, i.e., $t_1 \preceq t_2$ iff $\mathsf{won}(t_1, \pi) \leq \mathsf{won}(t_2, \pi)$. Then, since $\mathsf{won}(\cdot, \pi)$ corresponds to the score function where $p_{\mathsf{l}} = 0$ and $p_{\mathsf{w}} = 1$, we have that $\mathcal{T}$ is point-based.

Similarly to utility functions in Decision Theory [8], the following theorem shows that the score function associated to a point-based tournament is invariant under linear transformations.

THEOREM 2.2. *Let $\mathcal{T}$ be a point-based round-robin tournament and p_{l} and p_{w} be two real values such that $p_{\mathsf{l}} < p_{\mathsf{w}}$. Then, for each full path $\pi = \langle e_1, \ldots, e_h, \preceq\rangle$, we have that $a \preceq b$ iff* $\mathsf{score}(a, \pi) \leq \mathsf{score}(b, \pi)$.

PROOF. From $\mathcal{T}$ being point-based, there exist some $p'_{\mathsf{l}} < p'_{\mathsf{w}}$ such that, for each full path $\pi = \langle e_1, \ldots, e_h, \preceq\rangle$, the score function

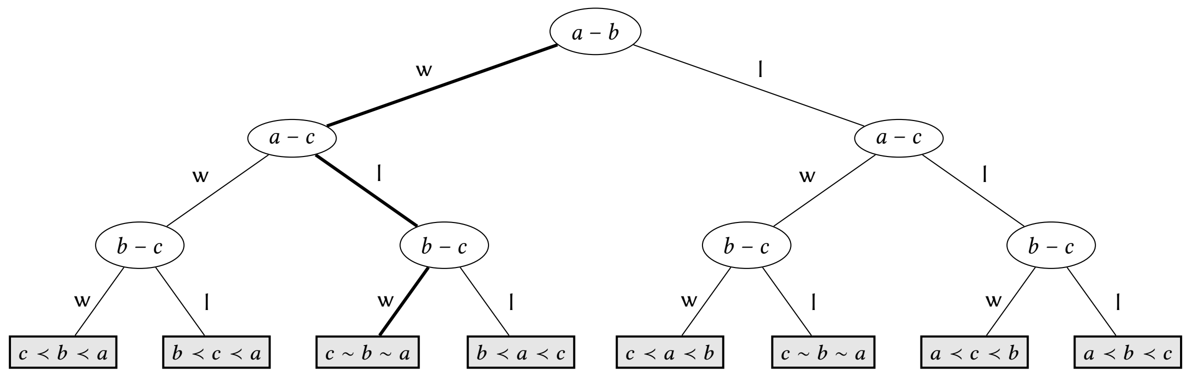

Figure 1: A tournament among three agents.

$\mathsf{score}'(\cdot,\pi) = \mathsf{won}(\cdot,\pi)p'_{\mathsf{w}} + \mathsf{lost}(\cdot,\pi)p'_{\mathsf{l}}$ is monotonic w.r.t. the ranking $\preceq$. By construction, we have that

$$\begin{aligned}\mathsf{score}(\cdot,\pi) &= \mathsf{won}(\cdot,\pi)p_{\mathsf{w}} + \mathsf{lost}(\cdot,\pi)p_{\mathsf{l}}\\ &= \mathsf{won}(\cdot,\pi)p_{\mathsf{w}} + (h - \mathsf{won}(\cdot,\pi))p_{\mathsf{l}}\\ &= \mathsf{won}(\cdot,\pi)(p_{\mathsf{w}} - p_{\mathsf{l}}) + hp_{\mathsf{l}}\,.\end{aligned}$$

Similarly, $\mathsf{score}'(\cdot,\pi) = \mathsf{won}(\cdot,\pi)(p'_{\mathsf{w}} - p'_{\mathsf{l}}) + hp'_{\mathsf{l}}$. Then, it is straightforward to see that $\mathsf{score}(\cdot,\pi) = \alpha\mathsf{score}'(\cdot,\pi) + \beta$, where

$$\alpha = \frac{p_{\mathsf{w}} - p_{\mathsf{l}}}{p'_{\mathsf{w}} - p'_{\mathsf{l}}} \quad \text{and} \quad \beta = p_{\mathsf{l}} - \frac{p_{\mathsf{w}} - p_{\mathsf{l}}}{p'_{\mathsf{w}} - p'_{\mathsf{l}}}\, p'_{\mathsf{l}}\,.$$

Since $\alpha > 0$, then $\mathsf{score}(\cdot,\pi)$ and $\mathsf{score}'(\cdot,\pi)$ are co-monotonic and hence the thesis. □

Roughly speaking, Theorem 2.2 says that the score function representing a point-based tournament does not depend on which weights $p_{\mathsf{l}} < p_{\mathsf{w}}$ we choose. Therefore, w.l.o.g. we will always assume that $p_{\mathsf{l}} = 0$ and $p_{\mathsf{w}} = 1$, that is, $\mathsf{score}(a,\pi) = \mathsf{won}(a,\pi)$.

Then, the *score vector* v_π of a full path π is the sequence of scores $\{\mathsf{won}(a,\pi)\}_{a\in A}$, ordered by non-decreasing value [9]. We say that a vector u of integers is a *round-robin score vector* if there exists a round-robin tournament $\mathcal{T}$ and a full path $\pi = \langle e_1,\ldots,e_h,\preceq\rangle$ in $\mathcal{T}$ such that u is the score vector of π. We recall the following classical result, called Landau's theorem.

Theorem 2.3 ([6]). *A vector $(u_1,\ldots,u_n) \in \mathbb{N}^n$ is a round-robin score vector if and only if, for all $k = 1,\ldots,n-1$,*

$$\sum_{i=1}^{k} u_i \geq \binom{k}{2} \qquad \textit{and} \qquad \sum_{i=1}^{n} u_i = \binom{n}{2}.$$

Next, define the following relation between paths: we say that two paths $\pi = \langle e_1,\ldots,e_h,\preceq\rangle$ and $\pi' = \langle e'_1,\ldots,e'_{h'},\preceq'\rangle$ are *homologous* if $(e_1,\ldots,e_h)$ is a permutation of $(e'_1,\ldots,e'_{h'})$. Consider a node x in a round-robin tournament and its children x_{w} and x_{l}. It is straightforward to see that the above relation induces a bijection between the paths starting from x_{w} and those starting from x_{l}.

The following lemma states that, once we fix the number of agents, all possible round-robin score vectors occur in all point-based round-robin tournaments. The proof is straightforward and left as an exercise to the reader.

Lemma 2.4. *Let $\mathcal{T}$ be a point-based round-robin tournament with n agents and let u be a round-robin score vector with n agents. Then, there exists a full path π in $\mathcal{T}$ such that $u = v_\pi$.*

3 NOTIONS OF RELEVANCE

In the following, unless differently specified, we assume a fixed number n of agents. Given two rankings $\preceq_1$ and $\preceq_2$ we say that an agent a prefers $\preceq_2$ to $\preceq_1$ if there are fewer agents above a in the former ranking or, those being equal, the number of ties is smaller. Formally, let $\mathsf{b}_a(\preceq) = |\{b \in A \mid a \prec b\}|$ be the number of agents that have a better placement than a in the ranking $\preceq$, and let $\mathsf{s}_a(\preceq) = |\{b \in A \mid a \sim b\}|$ be the number of agents having the same placement as a. We say that a prefers $\preceq_2$ to $\preceq_1$, denoted by $\mathcal{P}_a(\preceq_1,\preceq_2)$, iff:

(i) $\mathsf{b}_a(\preceq_2) < \mathsf{b}_a(\preceq_1)$, or
(ii) $\mathsf{b}_a(\preceq_2) = \mathsf{b}_a(\preceq_1)$ and $\mathsf{s}_a(\preceq_2) \leq \mathsf{s}_a(\preceq_1)$.

For instance, consider the two rankings $a \prec b \prec c$ and $a \sim b \sim c$. Agent c prefers the former, because it is the only agent on top. On the other hand, agent b prefers the latter, because there is no agent strictly above it.

The preference relation $\mathcal{P}_a$ is itself a weak order among rankings. We write $[\preceq]_a$ to denote the class of rankings $\preceq'$ that are a-equivalent to $\preceq$, i.e., $\mathcal{P}_a(\preceq,\preceq')$ and $\mathcal{P}_a(\preceq',\preceq)$. Since the number of possible rankings is finite, $\mathcal{P}_a$ induces a finite sequence $C^a_1,\ldots,C^a_m$ of a-equivalent classes, linearly ordered by $\mathcal{P}_a$.

We would like each match in a tournament to be *relevant* to both involved agents. Intuitively, a match is relevant for an agent if its future ranking depends on the outcome of that match. As anticipated in the Introduction, we adopt a probabilistic view: we assume that each agent is equipped with a *belief* about the likelihood of different outcomes for all future matches. We assume that agents are never absolutely certain about the outcome of future matches, so that the likelihood of a given outcome lies in the open interval $(0,1)$.

In the following, we fix a tournament $\mathcal{T}$ and omit it from the notation. Let $\mathcal{E}$ be the set of all events, a *belief* is a function $\beta : \mathcal{E} \to (0,1)$ such that for all agents a,b: *(i)* $\beta(a,b,\mathsf{w}) + \beta(a,b,\mathsf{l}) = 1$ and *(ii)* $\beta(a,b,\mathsf{w}) = \beta(b,a,\mathsf{l})$. A belief assigns probabilities to the outcomes of every possible match: the first condition derives from the fact that w and l are the only possible outcomes; the second condition simply states that matches are zero-sum games where if agent a wins then the opponent agent b loses, and vice versa.

For each internal node x, a belief naturally induces a distribution over all paths starting from x. The probability of any path $\pi = \langle e_1, \ldots, e_h, \preceq \rangle$ is simply obtained by combining the beliefs on each match, treated as independent events:

$$Pr_\beta(\pi) = \prod_{i=1}^{h} \beta(e_i)\,. \tag{1}$$

Note that homologous paths in a round-robin tournament have the same probability to occur according to any belief.

For an agent a, let R_a be the random variable assigning to each path $\pi = \langle e_1, \ldots, e_h, \preceq \rangle$ the a-equivalence class of the ranking $\preceq$, i.e., $R_a(\pi) = [\preceq]_a$. We denote by $\gamma(R_a, x, \beta)$ the distribution of R_a over all paths starting from x. More precisely, given an internal node x and an a-equivalence class C_j^a with $1 \leq j \leq m$, let $\mathsf{Paths}(x, a, j)$ be set of all paths starting from x whose final ranking belongs to C_j^a. Then, $\gamma(R_a, x, \beta)$ associates to the class C_j^a the probability

$$\sum_{\pi \in \mathsf{Paths}(x,a,j)} Pr_\beta(\pi).$$

Example 3.1. Consider again the example in Figure 1 and assume that a bookmaker estimates the following odds β: agent a has a probability 0.7 of beating b and a probability 0.5 of beating agent c; agent b has a probability 0.4 of beating c.

According to the final rankings, each agent t has four different equivalence classes $C_1^t, \ldots, C_4^t$, where C_1^t is the set of all rankings where t comes in last, C_2^t means that t is in the second position, C_3^t is the case where all the agents tie, and finally in C_4^t t takes the first position with no ties.

Thus, the class C_4^a corresponds in Figure 1 to the two leftmost paths. Let x be the root of the tree, $\gamma(R_a, x, \beta)$ assigns probability $0.7 * 0.5 = 0.35$ to C_4^a. Similarly, the paths where agent b takes the second position are the leftmost and the rightmost ones, hence C_2^b has probability 0.23.

Now, we formalize when agent a prefers an internal node x_1 to another node x_2, based on the corresponding distributions $\gamma_1 = \gamma(R_a, x_1, \beta)$ and $\gamma_2 = \gamma(R_a, x_2, \beta)$. This notion is based on 1^{st} order stochastic dominance between two distributions taking values over a linearly ordered set [4, 11]. In our case, the linear order is the sequence $C_1^a, \ldots, C_m^a$ of a-equivalence classes ordered by $\mathcal{P}_a$. Let Γ_1 and Γ_2 be the cumulative probabilities of γ_1 and γ_2, respectively:

$$\Gamma_i(C_k^a) = \sum_{j=1}^{k} \gamma_i(C_j^a)\,,$$

with $i = 1, 2$. Then, γ_2 *stochastically dominates* γ_1 if and only if, for all $1 \leq k \leq m$, $\Gamma_2(C_k^a) \leq \Gamma_1(C_k^a)$, with a strict inequality for at least one k.

Definition 3.2. A node x labelled with (a, b) is c-*relevant*, with $c \in \{a, b\}$, iff, for all beliefs β, $\gamma(R_c, x_{\mathsf{w}}, \beta)$ stochastically dominates $\gamma(R_c, x_{\mathsf{l}}, \beta)$ or vice versa. A tournament is *strongly* (resp., *weakly*) *relevant* if all internal nodes are relevant for both (resp., at least one of the) involved agents.

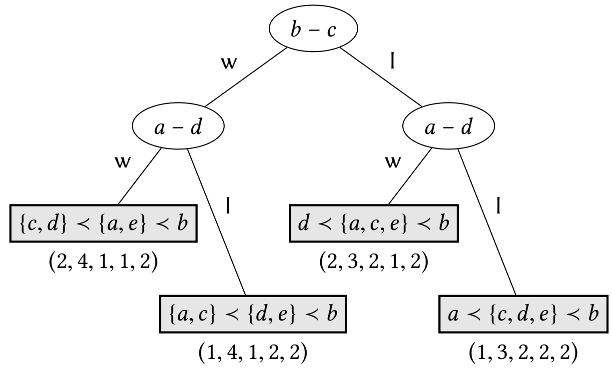

Figure 2: A fragment of a 5-agent tournament (see Example 3.3). Each leaf is decorated underneath by its score vector. Due to space constraints, agents that tie in a final ranking are grouped in set notation.

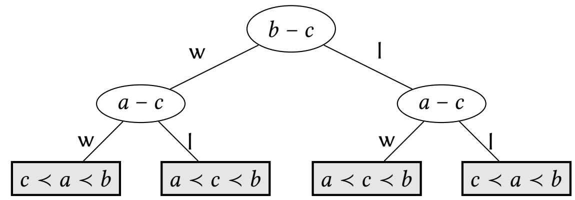

Figure 3: A fragment of a 3-agent tournament.

Example 3.3. Consider a static point-based round-robin tournament among five agents a, b, c, d, and e and let x be an internal node where the intermediate scores are:

a	b	c	d	e
1	3	1	1	2

and only two matches (b, c) and (a, d) are left. Figure 2 shows the subtree rooted in x.

Given a generic belief β, since at this point agent b will surely take first place with no ties, both distributions $\gamma(R_b, x_{\mathsf{w}}, \beta)$ and $\gamma(R_b, x_{\mathsf{l}}, \beta)$ assign probability 1 to the same equivalence class and zero to all the other classes. Consequently, none of the distributions stochastically dominates the other and hence x is b-irrelevant.

For round-robin tournaments, we now develop a characterization of relevance that avoids any reference to the probabilistic setting. For an agent t and an internal node x, we say that x is t-*important* if there exists a path from x_{w} whose ranking is not t-equivalent to the ranking of the homologous path starting from x_{l}. The following theorem states that, in round-robin tournaments, relevance implies importance.

Theorem 3.4. *For all nodes x in a round-robin tournament and all agents t, if x is t-relevant, then it is t-important.*

Proof. Assume that node x is t-relevant. Recall that for all beliefs β and all pairs of homologous paths π_1 and π_2, β assigns the same probability to them, i.e., $Pr_\beta(\pi_1) = Pr_\beta(\pi_2)$.

Assume by contradiction that x is not t-important. By definition, it follows that all pairs of homologous paths starting from

x_w and x_l, respectively, end with t-equivalent rankings. Hence, the two distributions $\gamma(R_t, x_w, \beta)$ and $\gamma(R_t, x_l, \beta)$ are equal, and so neither stochastically dominates the other, contradicting our initial assumption that x is t-relevant. □

Next, we prove that if the tournament is additionally point-based, relevance is equivalent to importance.

Theorem 3.5. *For all round-robin point-based tournaments, nodes x labeled with match (a, b), and agents $t \in \{a, b\}$, x is t-relevant iff it is t-important.*

Proof. Theorem 3.4 implies the "only if" direction, so we are left to prove the "if" implication. Assume that x is t-important, it follows that there exist two homologous paths $\pi_w = \langle e_1, \ldots, e_h, \preceq_w \rangle$ and $\pi_l = \langle e'_1, \ldots, e'_h, \preceq_l \rangle$ such that $\preceq_w$ and $\preceq_l$ are not t-equivalent. Assume w.l.o.g. that $t = a$. As the tournament is point-based, a strictly prefers $\preceq_w$ to $\preceq_l$, i.e., $\mathcal{P}_a(\preceq_l, \preceq_w)$ and not $\mathcal{P}_a(\preceq_w, \preceq_l)$.

Let $C_1^a, \ldots, C_m^a$ be the a-equivalence classes of rankings, ordered by $\mathcal{P}_a$. Let β be any belief and let Γ_w (resp., Γ_l) be the cumulative distribution of $\gamma(R_a, x_w, \beta)$ (resp., $\gamma(R_a, x_l, \beta)$). For all homologous paths (π_1, π_2) respectively starting from (x_w, x_l) and ending in rankings $(\preceq_1, \preceq_2)$, it holds $\mathcal{P}_a(\preceq_2, \preceq_1)$. Hence, if $\preceq_1$ belongs to the class C_i^a, $\preceq_2$ belongs to a class C_j^a with $j \geq i$. As a consequence, we have that $\Gamma_w(C_i^a) \leq \Gamma_l(C_i^a)$ for all $i = 1, \ldots, m$.

In order to prove stochastic dominance, it remains to show that there exists an i such that the above inequality is strict. Let C_k^a be the least preferred class in that linear order such that there exist two homologous paths (π_1, π_2) with π_2 ending in C_k^a and π_1 ending in a different, and hence strictly preferred, class. Such a class is guaranteed to exist, due to the pair of paths (π_w, π_l). Let $p = Pr_\beta(\pi_1) = Pr_\beta(\pi_2)$, we have that $p > 0$.

It follows that

$$\gamma(R_a, x_w, \beta)(C_i^a) = \gamma(R_a, x_l, \beta)(C_i^a)$$

for all $i = 1, \ldots, k-1$, and

$$\gamma(R_a, x_w, \beta)(C_k^a) \leq \gamma(R_a, x_l, \beta)(C_k^a) - p < \gamma(R_a, x_l, \beta)(C_k^a).$$

Hence, $\Gamma_w(C_k^a) < \Gamma_l(C_k^a)$, and our thesis. □

Example 3.6. Consider again Figure 2. Let x be the root labelled with (b, c) and $\preceq_1, \ldots, \preceq_4$ be the final rankings from left to right. Consider the two rankings $\preceq_1$ and $\preceq_3$. Since in the former c takes third place whereas in the latter c comes in second, we have that c strictly prefers $\preceq_3$ over $\preceq_1$, i.e., $\mathcal{P}_c(\preceq_1, \preceq_3)$ and not $\mathcal{P}_c(\preceq_3, \preceq_1)$. Now, those rankings appear at the end of the two homologous paths $\langle(a, d, \mathrm{w}), \preceq_1\rangle$ descending from x_w and $\langle(a, d, \mathrm{w}), \preceq_3\rangle$ descending from x_l. Consequently, x is c-important and, by Theorem 3.5, also c-relevant.

It is easy to see that being point-based is necessary for Theorem 3.5 to hold. In particular, consider Figure 3, which depicts a fragment of a 3-agent round-robin tournament. Notice that the tournament is not point-based. To see this, call $\preceq_1, \ldots, \preceq_4$ the four final rankings, left to right. In the path leading to $\preceq_3$, a (resp., c) has won one more match (resp., one fewer match) compared to the path leading to $\preceq_4$. This contradicts the fact that $a \prec_3 c$ and $c \prec_4 a$. Now, let x be the root of the fragment, and x_w and x_l be its left and right children, respectively. Observe that x is c-important but not c-relevant. Importance follows from comparing the two homologous paths starting from x_w and x_l and leading to $\preceq_l$ and $\preceq_3$, respectively. Then, consider the belief β that assigns equal probabilities to either agent winning the match (a, c). According to β, in both x_w and x_l agent c has probability $\frac{1}{2}$ of being second and probability $\frac{1}{2}$ of being third (i.e., $\gamma(R_c, x_w, \beta) = \gamma(R_c, x_l, \beta)$). So, neither distribution stochastically dominates the other and therefore x is not c-relevant.

4 EXISTENCE OF RELEVANT TOURNAMENTS

4.1 Static Tournaments

The following theorem states that, in every round-robin tournament for at least 5 agents that is static and point-based, there is a match that is irrelevant for both of the involved agents.

Theorem 4.1. *For all $n \geq 5$, there is no round-robin tournament that is static, point-based, and weakly relevant.*

Proof. Since the tournament is static, there is a match, say (i, j), that is played last in all paths. Consider a path, starting from the root, in which i wins all matches, j loses all matches, and the other $n-2$ agents win approximately half of their games, leading to a node x, where only the match (i, j) remains to be played. At x, agent i has score $n-2$, agent j has score 0, and the other agents have score approximately equal to $\frac{n-1}{2}$. Precisely, if n is odd then all other agents have score $\frac{n-1}{2}$, otherwise some of them have score $\lfloor \frac{n-1}{2} \rfloor$ and some have score $\lceil \frac{n-1}{2} \rceil$. It is easy to prove that this is indeed a valid scenario [2]. We prove that neither i nor j can change their ranking in their last game. Since $n \geq 5$, $\lfloor \frac{n-1}{2} \rfloor \geq 2$, so even if agent j wins the last match, its ranking will still be the the last of all. Dually, $\lceil \frac{n-1}{2} \rceil \leq n-3$, so even if agent i loses the last match, its ranking will be the first of all. Notice that the latter argument does not hold for $n \leq 4$. □

Example 4.2. For $n = 5$, the scenario described in Theorem 4.1 leads to the score vector $(s_1, \ldots, s_5) = (0, 2, 2, 2, 3)$, where the last match is between agents 1 and 5. Clearly, those agents cannot modify their ranking as a consequence of their last match.

The following proposition can be checked by inspection.

Proposition 4.3. *For all $n \leq 4$, there is a strongly relevant static round-robin tournament.*

4.2 Dynamic Tournaments

In this section, we describe a family $T(n)$ of dynamic round-robin point-based tournaments that are weakly relevant, for all numbers n of agents. In the following, for a score vector $(s_1, \ldots, s_n)$, we say that the vector $(\delta_0, \ldots, \delta_{n-1})$ is the corresponding *score difference vector*, where $\delta_0 = s_1$ and $\delta_i = s_{i+1} - s_i$ for all $i = 1, \ldots, n-1$.

Let $T(2)$ be the trivial round-robin tournament between 2 agents, $T(n+1)$ is recursively defined as follows. First, agents $1, \ldots, n$ play $T(n)$ (*first phase*), resulting in a score vector $(s_1, \ldots, s_n)$ and score differences $(\delta_0, \ldots, \delta_{n-1})$. We partition the agents into two sets:

- $\Delta_0 = \{i = 1, \ldots, n-1 \mid \delta_i = 0\}$;
- $\Delta_+ = \{i = 1, \ldots, n-1 \mid \delta_i \geq 1\} \cup \{n\}$.

[2]This scenario is related to the notions of *regular* and *nearly-regular tournament* in graph theory [9].

The *second phase* of the tournament consists in all matches involving agent $n+1$. First, we let $n+1$ play against all agents in Δ_+, in any order. Then, we schedule the remaining matches, in any order.

Example 4.4. Let $T(5)$ be a recursive tournament among the agents $\{a, b, c, d, e\}$ where the first phase consists in a tournament $T(4)$ among the first four agents. Then, consider as possible outcomes of the first phase the score vectors $u = (0, 2, 2, 2)$ and $v = (1, 1, 2, 2)$, where scores are assigned, from left to right, to a, b, c, and d. By construction, in the first case $\Delta_0 = \{b, c\}$ and $\Delta_+ = \{a, d\}$, whereas in the second case $\Delta_0 = \{a, c\}$ and $\Delta_+ = \{b, d\}$. Consequently, in the first case the match (a, e) will precede the match (b, e) whereas in the second case the opposite holds. This shows that different outcomes of $T(4)$ dynamically induce different schedules in the second phase. In particular, a possible schedule for the second case (score vector v) requires e to challenge the other agents in the following order: b, d, c, and a. If we used such a schedule in the first case (where $T(4)$ ends with the score vector u) and e beat b, d, and c, then the scores before the last match (e, a) would be:

a	b	c	d	e
0	2	2	2	3

and hence the last match would be irrelevant for both a and e.

The intuition behind the recursive tournament $T(n+1)$ is the following. Agents in Δ_0 are "safe" in the sense that, when they will play against $n+1$, the match will be relevant for them. The other agents are "unsafe" in the same sense. Hence, we first schedule the unsafe matches, because they are certainly relevant for $n+1$ (as a consequence of the next Lemma 4.5). Then, we schedule the remaining matches, that will be relevant for the other agent.

Lemma 4.5. *Assume that the sub-tournament $T(n)$ ends with score vector $(s_1, \dots, s_n)$ and maximum score difference $\delta^{\max}$. For all $i = 1, \dots, n$ such that $i \le s_n$ and $(n-i) \ge \delta^{\max}$, the $(i+1)$-th match of the second phase is relevant for agent $n+1$.*

Proof. Call x a node in the tournament corresponding to the $(i+1)$-th match of the second phase. At that point, agent $n+1$ has played i matches, so its current score $\hat{s}$ is at most i. Since $i \le s_n$, if we add a virtual agent 0 with score $s_0 = 0$, there is an index $j \in \{0, \dots, n-1\}$ such that $s_j \le \hat{s} \le s_{j+1}$. We now distinguish two cases.

First, assume $\hat{s} = s_j$. If agent $n+1$ loses all future matches, it ends up in a leaf y with score $\hat{s}$. On the other hand, if agent $n+1$ wins the next match and then loses all remaining matches, it ends up in a leaf z with score $\hat{s}+1$. One can prove that node z is always preferable, from the point of view of $n+1$, to node y, but the specific reason depends on the final score of agent j in y and z. In some cases, agents $n+1$ and j are tied in y, whereas $n+1$ is above j in z. In other cases, agent j is above agent $n+1$ in y, whereas they are tied in z. Notice that all other agents (except $n+1$) have the same score in y and z, except for a single agent which has one less point in z (another more reason for $n+1$ to prefer z).

In the following, we denote by s_k^y (resp., s_k^z) the score of agent k in node y (resp., z). For all agents $k \ne n+1$, it holds $s_k^y \ge s_k^z$. It follows that $\mathrm{b}_{n+1}(y) \ge \mathrm{b}_{n+1}(z)$. If $\mathrm{b}_{n+1}(y) > \mathrm{b}_{n+1}(z)$, by definition agent $n+1$ prefers z over y. Otherwise, it holds $\mathrm{b}_{n+1}(y) = \mathrm{b}_{n+1}(z)$ and we distinguish three sub-cases, based on the final score of agent j in y and z.

First sub-case: $s_j^y = s_j^z = s_j + 1$. Then, $s_{n+1}^y = \hat{s} = s_j$ whereas $s_{n+1}^z = \hat{s} + 1 = s_j + 1$. Hence, $j >^y n+1$ and $j \sim^y n+1$, which contradicts the assumption $\mathrm{b}_{n+1}(y) = \mathrm{b}_{n+1}(z)$.

Second sub-case: $s_j^y = s_j^z = s_j$. Then, $j \sim^y n+1$ and $j <^z n+1$. So, $\mathrm{s}_{n+1}(z) < \mathrm{s}_{n+1}(y)$ and we are done.

Third sub-case: $s_j^y = s_j + 1 > s_j^z = s_j$. Then, $j >^y n+1$ and $j <^z n+1$, which again contradicts $\mathrm{b}_{n+1}(y) = \mathrm{b}_{n+1}(z)$.

We are left to examine the case $\hat{s} > s_j$. Let $k = s_{j+1} + 1 - \hat{s}$. Under our assumptions, we have

$$k = s_{j+1} + 1 - \hat{s} \le s_{j+1} + 1 - (s_j + 1) = s_{j+1} - s_j \le \delta^{\max} \le n - i.$$

If $k = 0$, the argument is similar to the case $\hat{s} = s_j$, and compares the path where $n+1$ wins only the first match and the path where $n+1$ loses all matches.

If instead $k > 0$, we consider a path from x in which $n+1$ loses the first match and then wins exactly $k-1$ of the remaining matches, leading to a leaf y. As second path, consider the one in which $n+1$ wins the first match and then proceeds analogously to the first path, ending in a leaf z. In node y, agent $n+1$ achieves score $s_{n+1}^y = \hat{s} + k - 1 = s_{j+1}$, whereas in z we have $s_{n+1}^z = \hat{s} + k = s_{j+1} + 1$. We prove that agent $n+1$ prefers z over y. As before, for all agents $k \ne n+1$, it holds $s_k^y \ge s_k^z$. If $s_{j+1}^y = s_{j+1}$, then in y agents $n+1$ and $j+1$ have the same score s_{j+1}, whereas in z agent $n+1$ has one more point than $j+1$. Hence, z is preferable to y.

If instead $s_{j+1}^y = s_{j+1} + 1$, in y agent $j+1$ has one more point than $n+1$, whereas in z agent $j+1$ is either at the same score than $n+1$ (if $s_{j+1}^z = s_{j+1}^y$), or it has one less point than $n+1$ (if $s_{j+1}^z = s_{j+1} = s_{j+1}^y - 1$). □

The following lemma states that, whenever a score difference δ_i is higher than one, δ_i other score differences must be zero (a.k.a. *plateaus*).

Lemma 4.6. *For all score vectors and score difference vectors, for all i s.t. $\delta_i > 1$ it holds $|\Delta_0| \ge \delta_i$.*

Proof. Let $S_{j,k} = \sum_{l=j}^{k} s_l$. We have $S_{1,n} = S_{1,i} + S_{i+1,n} = \binom{n}{2}$. We can write $S_{i+1,n}$ as $s_i(n-i) + h$. By Landau's theorem, $S_{1,i} \ge \binom{i}{2}$. It follows that

$$h \le \binom{n}{2} - \binom{i}{2} - s_i(n-i). \tag{2}$$

We distinguish two cases.

First, assume that $s_i \ge i - 1$ (see Figure 4 for an illustration). By (2), $h \le \binom{n}{2} - \binom{i}{2} - (i-1)(n-i) = \binom{n-i+1}{2} = \sum_{j=1}^{n-i} j$. On the other hand, $h \le (n-i) + (n-i-1) + \ldots + 1$ and hence

$$(\delta_i - 1)(n-i) + \sum_{j=i+1}^{n-1} (\delta_j - 1)(n-j) \le 0. \tag{3}$$

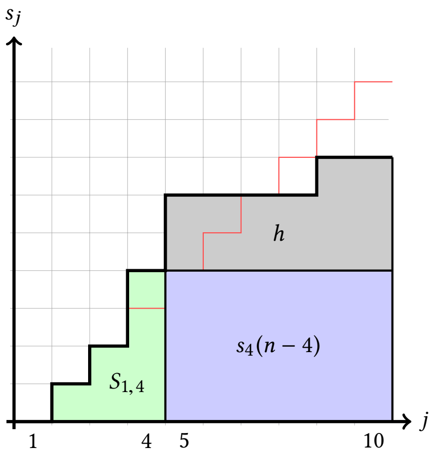

Figure 4: Case 1 of Lemma 4.6: $s_i \geq i-1$. The score vector is $(0,1,2,4,6,6,6,6,7,7)$, with $i=4$. We have $\delta_i=2$ and four plateaus, all occurring after position i.

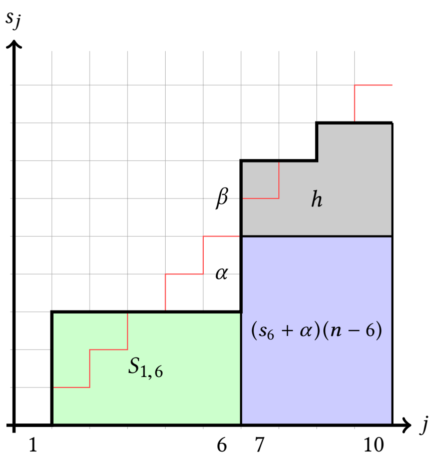

Figure 5: Case 2 of Lemma 4.6: $s_i < i-1$. The score vector is $(0,3,3,3,3,3,7,7,8,8)$, with $i=6$. We have $\delta_i=4$, four plateaus before position i and two more plateaus after it.

Assume by contradiction that $|\Delta_0| < \delta_i$, and let Δ_0' be the intersection between Δ_0 and the range $\{i+1,\ldots,n-1\}$. Then,

$$\begin{aligned}&(\delta_i-1)(n-i)+\sum_{j=i+1}^{n-1}(\delta_j-1)(n-j)\\ \geq&(\delta_i-1)(n-i)-\sum_{j\in\Delta_0'}(n-i-1)\\ \geq&(\delta_i-1)(n-i)-(\delta_i-1)(n-i-1)\\ =&\delta_i-1.\end{aligned}$$

From the above and (3), it follows $\delta_i-1\leq 0$, which contradicts the assumption $\delta_i>1$.

As second case, assume that $s_i < i-1$ and let $\alpha = i-1-s_i > 0$ (see Figure 5 for an illustration). First, we show that there are at least α plateaus before index i. Since $s_i = s_1 + \sum_{j=1}^{i-1}\delta_j$, and both s_1 and the δ_j's are non-negative integers, at most s_i different δ_j's are positive. The remaining $(i-1)-s_i$ must be zero, hence $|\Delta_0|\geq\alpha$. If $\delta_i\leq\alpha$, we are done. Otherwise, let $\beta=\delta_i-\alpha>0$. We can write $S_{i+1,n}$ as $(s_i+\alpha)(n-i)+h$. As in the previous case, $h\leq\binom{n-i+1}{2}$. On the other hand, $h=\beta(n-i)+\delta_{i+1}(n-i-1)+\ldots+\delta_{n-1}$. With an argument similar to the previous case, we obtain that the number of plateaus between $i+1$ and n is at least β. In conclusion, the total number of plateaus is at least $\alpha+\beta=\delta_i$. □

We can now prove the main result of this section.

Theorem 4.7. *The recursive tournament scheme $T(n+1)$ is weakly relevant, for all $n\geq 1$.*

Proof. Assume by induction that the sub-tournament $T(n)$ is weakly relevant.

First, we prove that the nodes in $T(n)$ are still weakly relevant when considered as nodes in the first phase of $T(n+1)$. Consider a node x in $T(n)$ labeled with (a,b) and assume that x is t-relevant in $T(n)$, for $t\in\{a,b\}$. By Theorem 3.5, x is t-important in $T(n)$. Hence, there is a path π_w from x_w in $T(n)$ whose ranking is not t-equivalent to the ranking of the homologous path π_l starting from x_l. Now, both π_w and π_l can be prolonged assuming that agent $n+1$ wins all matches in the second phase of $T(n+1)$. Notice that the two prolonged paths are homologous. Moreover, at the end of the two prolonged paths, all agents except $n+1$ have the same points that they had at the end of π_w and π_l, respectively. Therefore, x is t-important and weakly relevant in $T(n+1)$.

Next, we prove that the matches in the second phase are also weakly relevant. Assume that the sub-tournament $T(n)$ ends with score vector $(s_1,\ldots,s_n)$ and score differences $(\delta_0,\ldots,\delta_{n-1})$. Partition the agents into Δ_0 and Δ_+, as explained earlier. The second phase first performs all matches between agent $n+1$ and each agent in Δ_+.

Let $\hat{i}$ be the maximum i satisfying Lemma 4.5, we have that the first $\hat{i}+1$ matches in the second phase are relevant for agent $n+1$. Let $r=\hat{i}+1=1+\min\{s_n, n-\delta^{\max}\}$. Let $u=|\Delta_+|$ be the number of unsafe agents, we prove that $r\geq u$.

First case: $r=s_n+1$. Write s_n as $s_1+\sum_{i=1}^{n-1}\delta_i$. Notice that each agent j in Δ_+ has $\delta_j\geq 1$, so

$$u\leq\sum_{j\in\Delta_+}\delta_j\leq\sum_{i=1}^{n-1}\delta_i\leq s_n<r.$$

Second case: $r=n-\delta^{\max}+1$. By Lemma 4.6 applied to $\delta^{\max}$, we have that $|\Delta_0|\geq\delta^{\max}-1$. Hence, $u=n-|\Delta_0|\leq n-\delta^{\max}+1=r$ and we are done. It follows that all matches between $n+1$ and one of the agents in Δ_+ are relevant for $n+1$.

It remains to evaluate the matches between agent $n+1$ and the agents in Δ_0. Let x be a node corresponding to such a match, involving agents $n+1$ and $j\in\Delta_0$. Let x_l be the child of x where j has lost the match, and x_w the other child. Consider any path from x_l to a leaf y and consider its homologous path from x_w, leading to the leaf z. We prove that agent j prefers z over y. We have $s_j^y=s_j=s_{j+1}$ and $s_j^z=s_j+1=s_{j+1}+1$. Moreover, agent $j+1$ has the same score in y and z.

First, assume $s^y_{j+1} = s_{j+1}$ (i.e., agent $j+1$ lost against $n+1$). Then, agent j has the same ranking as $j+1$ in y, whereas j's ranking is above $j+1$ in z. Next, assume $s^y_{j+1} = s_{j+1} + 1$ (i.e., agent $j+1$ won against $n+1$). Then, agent j's ranking is below $j+1$ in y, whereas it is at the same level in z. In both cases, agent j prefers z to y. By definition, node x is j-important and by Theorem 3.5 it is j-relevant. In conclusion, the whole tournament is weakly relevant. □

It is an open problem whether there exist strongly relevant round-robin point-based tournaments for arbitrarily large n.

4.3 Balanced Tournaments

The tournament described in Section 4.2 is poorly balanced, in the sense that an agent may be put on hold for a long time, and then be asked to play all of its matches in a row. Clearly, this is undesirable for a real-world round-robin tournament, in which we would like all agents to play equally often, or at least approximately so.

Recall from Section 2 that a round-robin balanced tournament with n agents is organized into $n-1$ rounds, during which each agent plays a single match. This structure is reminiscent of tournaments that follow a regular, often weekly, schedule, with each agent playing against another every week. Unfortunately, the following result states that being balanced in this sense is incompatible with being strongly relevant.

Theorem 4.8. *For all even $n \geq 6$, there is no round-robin point-based balanced tournament that is strongly relevant.*

Proof. Let u be the score vector $(0, \frac{n}{2}, \frac{n}{2}, \ldots, \frac{n}{2})$. It is easy to prove that this is a valid score vector according to Theorem 2.3. Consider an arbitrary round-robin point-based balanced tournament $\mathcal{T}$. By Lemma 2.4, there is a full path π in $\mathcal{T}$ such that u is its score vector. Let i be the agent having score 0 at the end of π. Consider the last event involving agent i in π, and let x be the corresponding node. We know that i loses the match at x, because it will end up with score 0. We prove that the score vector at x is of the type $(0, \frac{n}{2} - 1, \ldots)$. Clearly, agent i has score 0 at x because scores are non-decreasing in time. Then, by definition of balanced tournament, x is one the last $\frac{n}{2}$ nodes in π. Consequently, between x and the end of π each agent plays at most one match. So, the score of each agent at x is either equal to its final score, or to its final score minus one. In particular, the agent that plays against i at x is one of the agents whose score increases by one. This implies that the lowest non-zero score at x is $\frac{n}{2} - 1$.

To prove that x is not i-relevant, it is sufficient to observe that, no matter whether i wins or loses at x, in all leaves i ends up with score at most 1, corresponding to minimum rank, with no ties. □

5 CONCLUSIONS

This paper introduces the notion of (ir)relevant match and provides some positive and negative results on certain classes of tournaments. In principle, positive results may be used in practice, to devise improved tournaments in which no irrelevant matches occur. Even in the scope of negative results, future refined analyses may be able to quantify the amount of irrelevant matches, and suggest ways to minimize them.

A number of open problems remain to be solved. On the presently studied class of round-robin tournaments, we do not know whether there exist strongly relevant dynamic tournaments for arbitrarily large numbers of agents, although we know that, if they exist, they cannot be balanced (Theorem 4.8).

Moreover, the present model can be extended in various directions. For instance, it would be interesting to accommodate more than two possible outcomes for each match, as sport matches often have multiple outcomes. The simplest addition would be the "tie" outcome. In that case, negative results based on the irrelevance of the last match, such as Theorem 4.1, should easily carry over to the extended model.

An orthogonal extension is to tournaments with *simultaneous matches*. In particular, a common occurrence in practice is the type of tournaments where every week all agents play against each other simultaneously. Interestingly, simultaneity induces two opposing consequences on match relevance. On the one hand, it restricts the number and type of schedules available, compared to a plain dynamic tournament. This effect can in principle hinder relevance, as it brings the model closer to a static tournament. On the other hand, in order to prove the relevance of a given match, we seek a future scenario in which the outcome of that match makes the difference. If two matches x and y are played simultaneously, when proving the relevance of x we can freely choose the outcome of y, and vice versa. This ability may facilitate relevance and does not apply to plain dynamic tournaments.

We did not specifically consider the class of *double* round-robin tournaments, in which every agent challenges each other twice. This class has special practical relevance, as it is employed in a vast majority of team sports. Some results can easily be extended to this class. For instance, it is easy to prove, along the lines of Theorem 4.1, that for at least three agents no static double round-robin tournament is strongly relevant.

REFERENCES

[1] R.A. Bradley and M.E. Terry. Rank analysis of incomplete block designs, I. The method of paired comparisons. *Biometrika*, 39, 1952.
[2] L. de Alfaro, V. Polychronopoulos, and N. Polyzotis. Efficient techniques for crowdsourced top-k lists. In *Proc. of IJCAI, Melbourne, Australia, August 19-25, 2017*, pages 4801–4805, 2017.
[3] S. Dobson and J. Goddard. Spectator demand for football. In *The economics of football*, chapter 11. Cambridge U.P., 2012.
[4] W. Russell J. Hadar. Rules for ordering uncertain prospects. *American Economic Review*, 59:25–34, 1969.
[5] N. Jennett. Attendances, uncertainty of outcome and policy in Scottish league football. *Scottish J. of Political Economy*, 31(2):176–198, 1984.
[6] H.G. Landau. On dominance relations and the structure of animal societies: III. The condition for a score structure. *Bull. Math. Biophys.*, 15:143–148, 1953.
[7] J.W. Moon. *Topics on tournaments.* Holtt, Rinehart and Winston, 1968.
[8] R.B. Myerson. *Game Theory: Analysis of Conflict.* Harvard University Press, 1997.
[9] K.B. Reid. Tournaments. In J.L. Gross, J. Yellen, and P. Zhang, editors, *Handbook of Graph Theory*. Chapman & Hall/CRC, 2nd edition, 2013.
[10] A. Rubinstein. Ranking the participants in a tournament. *SIAM J. Appl. Math.*, 38(1), 1980.
[11] L. Sauro. On the hierarchical nature of partial preferences over lotteries. *Autonomous Agents and Multi-Agent Systems*, 31(6):1467–1505, 2017.
[12] S. Szymanski. The economic design of sporting contests. In *The Comparative Economics of Sport*, pages 1–78. Palgrave Macmillan UK, London, 2010.
[13] F.L. Wauthier, M.I. Jordan, and N. Jojic. Efficient ranking from pairwise comparisons. In *Proc. of the 30th Int. Conf. on Machine Learning*, volume 28 of *PMLR*, pages 109–117, 2013.