THE UNIVERSITY OF WARWICK

## Manuscript version: Author's Accepted Manuscript

The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

## Persistent WRAP URL:

http://wrap.warwick.ac.uk/162489

## How to cite:

Please refer to published version for the most recent bibliographic citation information.

## Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

## Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk.

# Equilibrium Computation For Knockout Tournaments Played By Groups 

Grzegorz Lisowski, M. S. Ramanujan, and Paolo Turrini<br>University of Warwick<br>United Kingdom<br>\{Grzegorz.Lisowski,R.Maadapuzhi-Sridharan,P.Turrini\}@warwick.ac.uk


#### Abstract

In single-elimination knockout tournaments, participants face each other based on a starting seeding and progress to the next rounds by beating their direct opponents. In this paper we initiate the study of coalitional knockout tournaments, which generalise singleelimination knockout tournaments by allowing groups of players, or coalitions, to strategically select one of their members to take part in the tournament, following the starting seeding. We investigate the algorithmic properties of pure strategies Nash equilibria in these games under various setups, i.e., whether or not choices can be made at each round and whether or not tournament progression is important to the group. Despite the more complex tournament structure when compared to single-elimination, we provide (quasi-) polynomial-time algorithms for all cases. Our results can be applied to those tournaments where pre-play selection plays an important role, such as sport events or elections with run-off.


## KEYWORDS

Single-elimination Tournaments; Coalitional Strategies; Equilibria

## ACM Reference Format:

Grzegorz Lisowski, M. S. Ramanujan, and Paolo Turrini. 2022. Equilibrium Computation For Knockout Tournaments Played By Groups. In Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), Online, May 9-13, 2022, IFAAMAS, 9 pages

## 1 INTRODUCTION

Sport events are frequently organised as knockout tournaments among selected individuals or teams. In international competitions, for example, coaches choose from a nation's best sportspeople to maximise their country's winning chances, which requires studying the potential rivals at different stages of the tournament. Strategic selection is also relevant in those tournaments where progression matters, even when the chances of winning the entire tournament may be slim. For instance, clubs participating in the UEFA Champions League receive significant financial rewards for reaching each stage of the competition and so, selecting the right squad (which has to be before the beginning of the knockout stages) is critical.

Selection headaches of this type are of course not limited to sports. Many other competitive environments can be modelled as tournaments, where self-interested participants strive to make their best possible choice against the predicted choices of their potential opponents. If we consider presidential elections with run-off, for instance, we observe a similar phenomenon: parties put forward

[^0]one candidate each, with the top two in the first round competing against each other. Also in this case the right selection of a candidate heavily depends on the potential choices of the opposing parties.

Tournaments, understood as directed graphs over players (see, e.g. [23], [20], [12], [21]), were introduced in the game theory and social choice research, and have gained attention in computer science for their well-behaved computational properties (see e.g., [4]). In this paper we focus on what are possibly the simplest and bestknown tournaments, knockout (or single-elimination) tournaments, where participants are initially associated with the leaves of a full binary tree, and the winner of the match played between the players at a pair of sibling nodes proceeds to the next stage, i.e., the parent of these two nodes. Knockout tournaments have been extensively analysed in the computational social choice and algorithmic game theory communities, either in connection with social choice functions (see, e.g. [24], [19], [6]) or from the point of view of an external manipulator trying to rearrange the initial seeding to guarantee a certain player to be the winner (see, e.g. [14], [26], [16], [3]).

However, the problem of analysing the strategic behaviour of groups in knockout tournaments - be they tennis players or party candidates to choose from - has been largely overlooked.

Our contribution. In this paper we address this gap by extending standard (single-elimination) tournaments to account for the strategic behaviour of groups. Before the tournament starts, we allow each of them to make an independent choice of the best selection, or candidate, to put forward. We study the equilibrium behaviour of groups (or coalitions) in such tournaments from an algorithmic point of view. Our analysis spans three axes: A) whether coalitions choose their candidates once and for all - called oneshot tournaments (which we model as winner-take-all knockout tournaments) - or not - dynamic tournaments B) whether only winning matters - what we call win-lose tournaments - or also tournament progression - beyond win-lose; C) whether we focus on computing equilibria or on verifying a given one. Despite the complex tournament structure, we show polynomial-time or quasi-polynomial-time algorithms for all these cases (Table 1).

Related literature. Our results are connected to multiple research lines in the computational analysis of tournaments. In the context of social choice theory, attention has been paid to stable solutions [4] and to subclasses displaying desirable properties [5]. The existence of well-behaved solutions has also motivated the study of the complexity of their computation (as in e.g., [7], [8], [5]).

Our approach is closely related to strategic voting (see, e.g. [22]), the growing branch of social choice theory dealing with strategic behaviour in collective decision-making and, more specifically, to strategic candidacy [9]. In our case the electorate is composed of
coalitions who vote for one of their members, i.e., the one to put forward in the tournament, strategising over the outcome of the tournament - our social choice function - given the choices of the opponents. Along these lines, [17] recently studied, in the context of tournament solutions, how politicians can form a coalition so that a representative of their group is elected. Further, [11] conducted a complexity analysis of how parties can win elections based on the voters' preferences over the set of all potential candidates. In the context of voting with candidate selection on Hotelling-Downs spaces, equilibrium computation was shown to be not tractable in the general case [13].

Finally, our results are directly relevant to the problem of tournament manipulation, for example understanding how the seeding can be manipulated to force a particular winner, a problem extensively studied in the literature (by e.g. [27], [10], [2] [10], [26], [14], [16], [3], [18]) or whether a pair of players can reverse their comparison to make one of them the winner of the tournament (see, e.g., [1]).

Paper structure. In section 2 we provide the basic setup, notably the definitions of binomial arborescence, coalitional strategies and equilibrium solutions, together with some key observations. Section 3 analyses one-shot tournaments, providing algorithm for computing their equilibria, while Section 4 focuses on dynamic ones.

## 2 PRELIMINARIES

Tournaments. Let $[n]=\{1,2, \ldots, n\}$ be a set of players and let $D$ be a digraph ( $[n], E$ ) with $E \subseteq[n]^{2}$, representing the results of a round-robin tournament played between them. Here, for every pair of players, exactly one is picked to be the winner between the two, which is encoded in $D$ as an arc directed from the winner to the loser. So, for every pair of players $i, j$ with $i \neq j$, exactly one of $(i, j),(j, i)$ belongs to $E$. We shall call such a graph a tournament digraph. Also, if $(i, j) \in E$, we say that $i$ beats $j$.

Let $V(D)$ denote the set of vertices of $D$. A knockout tournament or single-elimination tournament (henceforth, SE-tournament) on $D$ is defined as a complete binary tree $T$ with $n$ leaves $L(T)$ and a bijective function $\pi: V(D) \rightarrow L(T)$ called the seeding, mapping the $n$ players to the $n$ leaves. Then, the winner of the knockout tournament corresponding to $\pi$ is determined recursively: the winner at a leaf $l$ is the player $j$ with $l=\pi(j)$, and the winner of the subtree rooted at a node $v$ is the winner of the match between the winners of the two subtournaments rooted at the children of $v$, as decided by the orientation of the unique arc between these two vertices. In this paper, we will only consider SE-tournaments based on perfect binary trees, implying that the number of players entering the tournament is a power of 2 . An example of a SE-tournament is depicted in the left side of Figure 1.

Arborescences. The proofs in the paper will often use the technical notion of binomial arborescence [26], which allows for a succinct formulation of the structural properties of SE-tournaments. An arborescence is a rooted directed tree such that all the arcs are directed away from the root.

Definition 1 ([26]). Let D be a tournament digraph. The set of binomial arborescences over $D$ is recursively defined as follows:

- Each $a \in V(D)$ is a binomial arborescence rooted at $a$;


Figure 1: On the left an SE-tournament with 4 players. On the right the corresponding binomial arborescence. Notice that the first-round matches are $\left(a_{1}, b_{2}\right)$ and $\left(d_{1}, c_{2}\right)$, and the only second-round match is ( $a_{1}, d_{1}$ ), which is won by $a_{1}$.


Figure 2: A coalitional structure of teams $A, B, C, D$, each with two members. In red, the arborescence-shaped selection.

- If, for somer $>0, T_{a}$ and $T_{b}$ are $2^{r-1}$-node binomial arborescences rooted at $a$ and $b$ respectively, then the tree $T$ resulting from adding an arc from $a$ to $b$ is the $2^{r}$-node binomial arborescence rooted at $a$.

An example of binomial arborescence is in the right of Figure 1. If a binomial arborescence $T$ is such that $V(T)=V(D)$, then $T$ is a spanning binomial arborescence (s.b.a.) of $D$. Intuitively, an s.b.a. can be used to compactly encode how an SE-tournament will evolve, following its tournament digraph.

As shown by [26], there is a formal connection between binomial arborescences and knockout tournaments.

Proposition 1 ([26]). Let $D$ be a tournament digraph and let $v^{\star} \in V(D)$. Then, there is a seeding of $V(D)$ such that the resulting knockout tournament is won by $v^{\star}$ iff $D$ has an s.b.a. rooted at $v^{\star}$.

As a result of this proposition, we shall interchangeably use the terms binomial arborescence and SE-tournament, when these are clear from the context. Henceforth, we will mainly work with binomial arborescences, as this allows for neater proofs and procedures.

Subtournaments. Consider now a set of players $\left[2^{m}\right]$. We denote by $S E_{\pi,\left[2^{m}\right]}$ the spanning binomial arborescence representing the SE-tournament played by the players in $\left[2^{m}\right]$ following the seeding $\pi$. We call the root of $S E_{\pi,\left[2^{m}\right]}$ the winner of $S E_{\pi,\left[2^{m}\right] \text {. Moreover, }}$ for each $r \in\{0, \ldots, m\}$, we denote by $S E_{\pi,\left[2^{m}\right]}^{r}$ the binomial arborescence representing the subtournament of $S E_{\pi,\left[2^{m}\right]}$ played by the winners of all the $r^{\text {th }}$ round matches of $S E_{\pi,\left[2^{m}\right]}$ (notice that there are exactly $2^{m-r}$ players who win at least $r$ rounds). In other words, $S E_{\pi,\left[2^{m}\right]}^{0}$ is the same as the full tournament $S E_{\pi,\left[2^{m}\right]}$ and, for every $r \in[m], S E_{\pi,\left[2^{m}\right]}^{r}$ denotes the subgraph of $S E_{\pi,\left[2^{m}\right]}^{r-1}$ obtained by simply deleting all its leaves. Notice that $S E_{\pi,\left[2^{m}\right]}^{m}$ contains a single vertex which corresponds to the winner of $S E_{\pi,\left[2^{m}\right]}$.

Furthermore, let $u, v_{1}, v_{2}$ be vertices of a binomial arborescence $T$ such that $u$ is the parent of $v_{1}$ and $v_{2}$. For each $x \in\{1,2\}$, let
$\ell_{x}$ denote the size (number of vertices) of the largest binomial arborescence contained in $T$ and rooted at $v_{i}$. If $\ell_{1}<\ell_{2}$, then we say that $v_{2}$ is a heavier child of $u$ than $v_{1}$. We denote this by $v_{2}>_{u}^{T} v_{1}$, with the explicit references to $u$ and $T$ dropped when these are clear from the context. Notice that $\ell_{1}$ and $\ell_{2}$ can never be equal. If $v_{2}$ is the unique child of $u$ that is heavier than all of their siblings, then we say that $v_{2}$ is the heaviest child of $u$. The lightest child of $u$ is symmetrically defined. For $r \in[\log \ell]$, these notions can be extended in the natural way to $r$ heaviest children (or $r$ lightest children) of $u$. Intuitively, assuming $u$ is not a leaf, the heaviest child of $u$ is the last player beaten by $u$ in the knockout tournament and the lightest child of $u$ is the first player beaten by $u$, i.e., the opponent of $u$ in the first round. Similarly, the $r$ heaviest children of $u$ are the last $r$ opponents of $u$ that lose to $u$ and the $r$ lightest children of $u$ are the first $r$ opponents of $u$ that lose to $u$.

Coalitional Structures. In this paper we study coalitional tournaments, where players are partitioned into groups (also called teams or coalitions). Informally, each coalition will put forward one of their members in the spot allocated to them by the seeding. When each coalition makes their choice, this will give rise to a standard SE-tournament. So, for a set $\mathcal{P}$ of players, a coalitional structure of $\mathcal{P}$ is a partition $C=\left\{C_{1}, C_{2}, \ldots, C_{\ell}\right\}$ of $\mathcal{P}$, such that $\ell=2^{m}$ for some $m>0$. For a set $C^{\prime} \subseteq C$ of coalitions, we use $\mathcal{P}\left(C^{\prime}\right)$ to denote their set of players. That is, $\mathcal{P}\left(C^{\prime}\right)=\bigcup_{C \in C^{\prime}} C$. Moreover, for a player $p \in \mathcal{P}$, we denote by $\mathcal{C}(p)$ the coalition to which $p$ belongs.

We use $D_{C}$ to denote the $|C|$-partite coalitional structure representing the results of matches played between all pairs of players of distinct coalitions. So, $D_{C}$ is a graph $(\mathcal{P}, E)$ such that for every pair of players $p_{i}, p_{j} \in \mathcal{P}$ for which $p_{i} \in C_{i}, p_{j} \in C_{j}$ and $C_{i} \neq C_{j}$, exactly one of $\left(p_{i}, p_{j}\right),\left(p_{j}, p_{i}\right)$ belongs to $E$.

Figure 2 shows a 4-partite coalitional structure, with a strategic selection reflecting the SE-tournament and the binomial arborescence in Figure 1.

Input representation. In order to analyze the running time of our algorithms, we need to first establish the input representation and input size. Given the observations above, we represent a coalitional knockout tournament in space $O\left(n^{2} \log n\right)+O(n \log \ell)+O(\ell \log \ell)$, where $\ell$ is the number of coalitions and $n$ the total number of players. Here, $O\left(n^{2} \log n\right)$ is a bound on the representation of $D_{C}$, $O(n \log \ell)$ is the space to represent the partition of players into coalitions and $O(\ell \log \ell)$ the space to represent the seeding. As the number of coalitions $\ell$ is at most the number of players $n$ (we do not permit empty coalitions), the input has bit-size $O\left(n^{2} \log n\right)$.

We assume that the seeding is fixed and known to all coalitions a priori. Hence, when possible, in order to keep the notation simple, we will refrain from explicitly referring to the seeding.

## 3 ONE-SHOT KNOCKOUT TOURNAMENTS

In this section, we study tournaments where the coalitions choose a representative once and for all, before the competition starts. We are interested in equilibrium strategies - specifically, Nash equilibria - and to study the complexity of their computation and verification.

Starting from a coalition structure $C$ we call a selection of players by all teams a strategy profile (or simply, a profile) and formally define it as a function $s: C \rightarrow \mathcal{P}$, such that $s\left(C_{i}\right) \in C_{i}$. In other words,
a strategy profile is a selection by each coalition of one of their members, an example of which is shown in Figure 2. We note that a strategy profile $s$ can be simply rewritten as a tuple ( $p_{1}, \ldots, p_{\ell}$ ) of members of the corresponding coalitions. Also, whenever the seeding $\pi$ is clear from the context, we write $S E_{S}$ to denote the tournament between players in $s$ following $\pi$, and $S E_{s\left[p_{i}^{\prime}\right]}$ the tournament obtained from $S E_{S}$ by replacing player $p_{i}$ with $p_{i}^{\prime}$ (in the same coalition as $p_{i}$ ). Furthermore, we say that a coalition $C$ wins tournament $S E_{s}$ if the winner of $S E_{S}$ belongs to $C$. Figure 1 features the tournament $S E_{\left\{a_{1}, b_{2}, c_{2}, d_{1}\right\}}$, with $a_{1}$ being the winner.

### 3.1 Win-Lose Games

The simplest type of one-shot tournaments we look at are the ones where only winning matters (i.e., win-lose games). When picking a player to put forward in a win-lose game, a coalition will have to reason about the possible outcomes of the tournament as a function of the opponents' choices. To encode this reasoning we define appropriate variants of the classical Nash equilibria.

Definition 2 (Equilibria). A profiles $=\left(p_{1}, \ldots, p_{\ell}\right)$ is $a$ Nash Equilibrium (NE), if for all $i$ and for all $p_{i}^{\prime} \in C_{i}:$ if $C_{i}$ wins $S E_{s\left[p_{i}^{\prime}\right]}$ then $C_{i}$ wins $S E_{s}$.

In other words, in a NE, no coalition can improve their outcome by changing their representative unilaterally. Clearly, a NE need not exist in general, which opens important algorithmic questions regarding their existence and computation.

### 3.1.1 Structural and algorithmic properties of NE.

Equilibrium existence. We start by addressing the question of deciding whether a given strategy profile is a NE, what we call the problem of recognising a NE. To do so we introduce a technical notion and a key lemma, before we are able to show its tractability.

Definition 3. Let $j \in[\ell]$ such that $p_{j}$ is the root of the s.b.a. $S E_{s}$. Let $j^{\prime}$ be a coalition distinct from $j$ and consider the $p_{j}-p_{j^{\prime}}$ path $H$ in $S E_{s}$. For every player $p$ on this path, we denote by Opponent ${ }_{H}^{+}\left[p, p_{j^{\prime}}\right]$ the set defined as follows.

- If $p=p_{j^{\prime}}$, Opponent ${ }_{H}^{+}\left[p, p_{j^{\prime}}\right]$ is the set of children of $p_{j^{\prime}}$ in $S E_{s}$.
- Else, Opponent ${ }_{H}^{+}\left[p, p_{j^{\prime}}\right]$ is the set $\{p\} \cup\left\{v \mid v\right.$ is a child of $p$ and $v>_{p}$
$\left.p^{\prime}\right\}$, where $p^{\prime}$ denotes the unique child of $p$ contained in $H$.
The utility of the above definition is derived from the fact that it formalises the set of future opponents that the coalition $C_{j^{\prime}}$ would have to face if it were to replace $p_{j^{\prime}}$ with a different player $p_{j^{\prime}}^{\prime}$, in the profile $s$ - what we denote $\bigcup_{p \in V(H)}$ Opponent ${ }_{H}^{+}\left[p, p_{j^{\prime}}\right]$. Note that this set is no larger than $\log \ell$ since that is the maximum number of opponents faced by any player in the tournament. We drop the explicit reference to $H$ in this notation when the root $p_{j}$ (the winner) is clear from the context. See Figure 3 for a visual illustration.

The following structural lemma forms the crux of our algorithmic results. Informally, we use the fact that a profile is a NE if and only if no coalition which is losing the tournament can switch their representative to one which beats all players beaten by the original representative plus every future opponent (that the original representative did not get to face following its loss), along with the properties of the technical notion in Definition 3.


Figure 3: An example of a non-NE strategy, where players 12 and 33 belong to the same coalition (for convenience called Coalition 12), and all other coalitions consist of one player only and the index of the Coalition is the same as the index of the player (e.g., Player 14 plays for Coalition 14). There are 33 players and 32 coalitions in total. The seeding pairs up Coalition $2 i+1$ and Coalition $2 i+2$ in the first round, for each $i \in\{0, \ldots, 15\}$. Consider the path $1 \rightarrow 9 \rightarrow 11 \rightarrow 12$, call it $H$. Notice that Opponent ${ }^{+}[12,12]=\emptyset$, Opponent ${ }^{+}[11,12]=$ $\{11\}$, Opponent $^{+}[9,12]=\{9,13\}$, Opponent $^{+}[1,12]=\{1,17\}$. Thus, if Coalition 12 wanted to win by replacing its chosen player (Player 12) with an alternative, then it would have to beat the players chosen by Coalitions $\{11,9,13,1,17\}$, which are precisely the Coalitions indexed by $\bigcup_{p \in V(H)}$ Opponent $_{H}^{+}[p, 12]$. Also, fixing the choices of all coalitions, player 33 improves upon player 12 , beating all the potential future opponents - Players $\{11,9,13,1,17\}$.

Lemma 1. Consider a profile $s=\left(p_{1}, \ldots, p_{\ell}\right)$ and let $j \in[\ell]$ such that $p_{j}$ is the root of the s.b.a. $S E_{s}$. Then, s is a NE iff there is no $j^{\prime} \in[\ell] \backslash\{j\}$ and player $p_{j^{\prime}}^{\prime} \in C_{j^{\prime}}$ such that $p_{j^{\prime}}^{\prime}$ beats all players in Opponent ${ }^{+}\left[p, p_{j^{\prime}}\right]$ for every player $p$ on the $p_{j}-p_{j^{\prime}}$ path in $S E_{s}$.

Figure 3 shows a non-NE profile based on the characterization in Lemma 1. As a consequence of this lemma, we obtain a simple algorithm for testing whether a given profile is a NE.

## Proposition 2. Recognising a NE is P-time solvable.

We note that this test can be done in $O\left(n^{3}\right)$-time, implying a running time that is subquadratic in the input size.

Equilibrium computation. We are further interested in computing a NE, which is much more complex than recognising one. Surprisingly, we show that this can still be done in quasi-polynomial time.

To be precise, our main theorem, Theorem 1, will show the existence of an $n^{O(\log \ell)}$-time algorithm for computing a NE. To obtain this result, we will make use of a key lemma, Lemma 2, which establishes that a NE (if one exists) can be obtained by composing specific types of strategies for various subtournaments. This lemma effectively implies that we can compute a NE (if one exists) by examining only at most $n^{O(\log \ell)}$ out of the set of possibly $n^{\Omega(\ell)}$ _ many strategy profiles.

Let us start with providing several useful notions. For a set $S$ and $q \in \mathbb{N}$, we denote by $\binom{S}{\leq q}$ the set of subsets of $S$ of size at most $q$ and by $[q]_{\geq 0}$ the set $\{0\} \cup[q]$. Let us denote by $C_{r}\left(C_{j}\right)$ the set of coalitions who could potentially meet Coalition $C_{j}$ within the first $r$-rounds (over all possible digraphs $D_{C}$ and profiles $s$ ) plus the Coalition $C_{j}$ itself. Notice that $\left|C_{r}\left(C_{j}\right)\right|=2^{r}$. Thus, $C_{r}\left(C_{j}\right) \backslash$ $C_{r-1}\left(C_{j}\right)$, denotes the set of all possible opponent coalitions that $C_{j}$ could face exactly in the $r^{\text {th }}$ round. For technical reasons, we set $C_{r}\left(C_{j}\right)=\emptyset$ for every $r \in\{-1,-2\}$ and $C_{j} \in C$. Finally, for the seeding $\pi$ we define by $\pi_{j}^{r}$ the restriction of $\pi$ to $C_{r}\left(C_{j}\right)$. Similarly, for a profile $s$, we denote by $s_{j}^{r}$, the profile $\left.s\right|_{C_{r}\left(C_{j}\right)}$, which is the restriction of $s$ to $C_{r}\left(C_{j}\right)$. Conversely, we say that $s$ is an extension of $s_{j}^{r}$. Notice that there can be multiple possible extensions of $s_{j}^{r}$.

We next define a mapping $\zeta$ which allow us to reason about a NE in terms of the outcomes of specific subtournaments.

Definition 4. We define the function $\zeta: \mathcal{C} \times \mathcal{P} \times \mathcal{P} \times[\log \ell]_{\geq 0} \times$ $\underset{\binom{\mathcal{P}}{\leq \log \ell}}{\underset{2}{S}} \mathbf{2}^{\mathcal{S}}$ as follows. Let $C_{j} \in C, r \in[\log \ell]_{\geq 0}, p_{j} \in C_{j}$, $p_{j^{\prime}} \in C_{j^{\prime}}$ such that $C_{j^{\prime}} \in C_{r}\left(C_{j}\right) \backslash C_{r-1}\left(C_{j}\right), Z \in(\underset{(\log \ell}{\mathcal{P}})$ such that $Z \cap \mathcal{P}\left(C_{r}\left(C_{j}\right)\right)=\left\{p_{j}\right\}$. Then, $\zeta\left(C_{j}, p_{j}, p_{j^{\prime}}, r, Z\right)$ denotes the set of all profiles s over $C_{r}\left(C_{j}\right)$ such that:

- (a) player $p_{j} \in C_{j}$ wins $S E_{\pi_{j}^{r}, s}$
- (b) player $p_{j^{\prime}}$ is the final opponent of $p_{j}$ in tournament $S E_{\pi_{j}}$,s
- (c) for every coalition $C_{j^{\prime \prime}} \in C_{r}\left(C_{j}\right) \backslash\left\{C_{j}\right\}$ and player $p_{j^{\prime \prime}} \in$ $C_{j^{\prime \prime}}$, either (i) $p_{j^{\prime \prime}}$ is beaten by a player in $Z$ or (ii) by a player in the set Opponent ${ }^{+}\left[p, p_{j^{\prime \prime}}\right]$ for some player $p$ on the $p_{j}-p_{j^{\prime \prime}}$ path in $S E_{\pi_{j}^{r}, s}$.
For every other choice of $C_{j}, p_{j}, p_{j^{\prime}}, r, Z, \zeta\left(C_{j}, p_{j}, p_{j^{\prime}}, r, Z\right)=\emptyset$.
So, $\zeta\left(C_{j}, p_{j}, p_{j^{\prime}}, r, Z\right)$ is the set of all those profiles over the $r$ round tournament comprised of Coalition $C_{j}$ 's first $r$ matches, such that the player $p_{j}$ (from Coalition $C_{j}$ ) wins these $r$ rounds, Player $p_{j}^{\prime}$ is the final opponent of $p_{j}$ in the $r$ round tournament (i.e., $p_{j}$ and $p_{j}^{\prime}$ play in the $r^{\text {th }}$ round) and all players in any coalition who could potentially meet $C_{j}$ within the first $r$ rounds are beaten either by a player in the set $Z$ or by a player in the set Opponent ${ }^{+}\left[p, p_{j^{\prime \prime}}\right]$ for some player $p$ on the $p_{j}-p_{j^{\prime \prime}}$ path in the s.b.a $S E_{\pi_{j}^{r}, s}$.

An informal description of the motivation behind this definition is the following. If we take $Z$ to denote $p_{j}$ plus its "potential future opponents" from round $r+1$ onwards, then $\zeta\left(C_{j}, p_{j}, p_{j^{\prime}}, r, Z\right)$ contains precisely all those profiles over $C_{r}\left(C_{j}\right)$ such that in any profile (over all coalitions) that is an extension thereof, there is no benefit for any coalition in $C_{r}\left(C_{j}\right)$ (other than possibly $C_{j}$ ) to unilaterally alter its strategy because any alternative player will either fail to win the first $r$ rounds anyway (Property (c) (ii) of Definition 4) or eventually lose to a "future opponent" in the set $Z$ (Property (c) (i) of Definition 4). Notice that the set $Z$ is never larger than $\log \ell$ because the number of rounds in the whole tournament is $\log \ell$. Our quasi-polynomial running time arises from the fact that we have $n^{O(\log \ell)}$ possibilities for $Z$ and essentially go over all of these possibilities in our algorithm.

Definition 4 and Lemma 1 entail that by setting the arguments of the function $\zeta$ appropriately, one can capture all NEs. Thus:

Observation 1. Every NE is contained in $\zeta\left(C_{j}, p_{j}, p_{j^{\prime}}, \log \ell,\left\{p_{j}\right\}\right)$ for some $C_{j} \in C, p_{j} \in C_{j}$ and $p_{j^{\prime}} \in C_{j^{\prime}}$ and conversely, for every $C_{j} \in$
$\mathcal{C}, p_{j} \in C_{j}$ and $p_{j^{\prime}} \in C_{j^{\prime}}$, every profile in $\zeta\left(C_{j}, p_{j}, p_{j^{\prime}}, \log \ell,\left\{p_{j}\right\}\right)$ is a NE.

Given the above observation, what naturally suffices is to compute the function $\zeta$ for all possible settings of the arguments. However, this may not be possible in quasi-polynomial time because even just listing all possible NEs could be too computationally expensive. To overcome this obstacle, we next prove a structural result that shows, intuitively, that if a strategy profile $s$ is a NE for a given tournament, then we can reconstruct $s$ (or an alternate NE) by going over the subtournaments at every possible level, examining the image (which is a set of profiles) of $\zeta$ by setting the arguments appropriately, and then merging these profiles. In particular, we show that if there is a NE-strategy $s$ then, for every round $r \in[\log \ell], s$ extends some profile $s^{\prime}$ in $\zeta\left(C_{j}, p_{j}, p_{j^{\prime}}, r, Z\right)$ for some $C_{j}, p_{j} \in C_{j}, p_{j^{\prime}} \in C_{j^{\prime}}$ and $Z$, and moreover, for any other profile $s^{\prime \prime} \in \zeta\left(C_{j}, p_{j}, p_{j^{\prime}}, r, Z\right)$, we can "cut" $s^{\prime}$ from $s$ and "paste" $s^{\prime \prime}$ to obtain another NE. The main algorithmic consequence of this fact is that for any choice of the arguments $C_{j}, p_{j}, p_{j^{\prime}}, r$ and $Z$, instead of computing all the profiles contained in $\zeta\left(C_{j}, p_{j}, p_{j^{\prime}}, r, Z\right)$, it is sufficient to compute a single one because if any of them can be extended to a NE, then every one of them can be extended to a NE. This "pruning lemma" is at the heart of our algorithm.

Lemma 2. Let $s=\left(p_{1}, \ldots, p_{\ell}\right)$ be a $N E, r \in[\log \ell]$ and let the winner of $S E_{s}$ be $p_{j^{\star}} \in C_{j^{\star} \text {. }}$. Let $p_{j} \in C_{j}$ be the winner of the binomial sub-arborescence $S E_{\pi_{j}^{r}}, s_{j}^{r}$ and let $p_{j^{\prime}} \in C_{j^{\prime}}$ denote the final opponent of $p_{j}$ in $S E_{\pi_{j}^{r}, s_{j}^{r}}$. Let $Z=\bigcup_{p \in V(H) \backslash\left\{p_{j^{\prime}}\right\}}$ Opponent ${ }_{H}^{+}\left[p, p_{j^{\prime}}\right]$ where $H$ denotes the $p_{j^{\star}-p_{j}{ }^{\prime}}$ path in $S E_{s}$.
(1) Then, $\zeta\left(C_{j}, p_{j}, p_{j^{\prime}}, r, Z\right)$ is non-empty.
(2) Moreover, for every $p_{j_{1}} \in C_{j_{1}}$ and $p_{j_{2}} \in C_{j_{2}}$, for every $\hat{s}_{1} \in$ $\zeta\left(C_{j}, p_{j}, p_{j_{1}}, r-1, Z \cup\left\{p_{j^{\prime}}\right\}\right)$, and $\hat{s}_{2} \in \zeta\left(C_{j^{\prime}}, p_{j^{\prime}}, p_{j_{2}}, r-1, Z \cup\right.$ $\left.\left\{p_{j^{\prime}}\right\}\right)$, the composed profile $\hat{s}_{1} \cdot \hat{s}_{2}$ is contained in $\zeta\left(C_{j}, p_{j}, p_{j^{\prime}}, r, Z\right)$.

Proof. We first show that $s_{j}^{r} \in \zeta\left(C_{j}, p_{j}, p_{j^{\prime}}, r, Z\right)$. By the definition of $s_{j}^{r}, p_{j}$ is the root of $S E_{\pi_{j}^{r}, s_{j}^{r}}$ and $p_{j^{\prime}}$ is the heaviest child of $p_{j}$ in $S E_{\pi_{j}^{r}, s_{j}^{r}}$. Hence, if $s_{j}^{r} \notin \zeta\left(C_{j}, p_{j}, p_{j^{\prime}}, r, Z\right)$, then there exists $p_{j^{\prime \prime}} \in C_{j^{\prime \prime}} \in C_{r}\left(C_{j}\right) \backslash\left\{C_{j}\right\}$, such that $p_{j^{\prime \prime}}$ beats every player in $Z$ and $p_{j^{\prime \prime}}$ beats every player in the set Opponent ${ }^{+}\left[p, p_{j^{\prime \prime}}\right]$ for every player $p$ on the $p_{j}-p_{j^{\prime \prime}}$ path in $S E_{\pi_{j}^{r}, s_{j}^{r}}$. Along with our choice of $Z$, this would then also imply that $p_{j^{\prime \prime}}$ beats every player in the set Opponent ${ }^{+}\left[p, p_{j^{\prime \prime}}\right]$ for every player $p$ on the $p_{j^{\star}}-p_{j^{\prime \prime}}$ path in $S E_{s}$, contradicting our choice of $s$ as a NE (see Lemma 1).

For the second statement, we need to prove that Properties (a), (b), and (c) in Definition 4 are satisfied by $\hat{s}=\hat{s}_{1} \cdot \hat{s}_{2}$. The first two properties are satisfied since $p_{j}$ beats $p_{j^{\prime}}, p_{j}=S E_{\hat{s}_{1}}^{r-1}$, and $p_{j^{\prime}}=S E_{\hat{s}_{2}}^{r-1}$. Suppose Property (c) is violated. Then, for some $p_{j^{\prime \prime}} \in C_{j^{\prime \prime}} \in \mathcal{C}_{r}\left(C_{j}\right) \backslash\left\{C_{j}\right\}, p_{j^{\prime \prime}}$ beats all players in $Z$ and all players in the set Opponent ${ }^{+}\left[p, p_{j^{\prime \prime}}\right]$ for each player $p$ on the $p_{j}-p_{j^{\prime \prime}}$ path in $S E_{\pi_{j}^{r}, \hat{,}}$. However, note that if $C_{j^{\prime \prime}} \in C_{r-1}\left(C_{j}\right)$ then we obtain a contradiction to our assumption that $\hat{s}_{1} \in \zeta\left(C_{j}, p_{j}, p_{j_{1}}, r-\right.$ $\left.1, Z \cup\left\{p_{j^{\prime}}\right\}\right)$ satisfies Property (c) of Definition 4, and otherwise (if $C_{j^{\prime \prime}} \in C_{r}\left(C_{j}\right) \backslash C_{r-1}\left(C_{j}\right)$ ), this violates our assumption that $\hat{s}_{2} \in \zeta\left(C_{j^{\prime}}, p_{j^{\prime}}, p_{j_{2}}, r-1, Z \cup\left\{p_{j^{\prime}}\right\}\right)$ satisfies Property (c) of Definition 4.

Lemma 2 forms the core of our algorithm to compute a NE, whose running time we establish with the following theorem.

Theorem 1. There is an $n^{O(\log \ell)}$-time algorithm for computing $a$ NE.

Proof. Due to Observation 1, our algorithm to compute a NE aims to identify and return an element of a non-empty $\zeta\left(C_{j}, p_{j}, p_{j^{\prime}}\right.$, $\left.\log \ell,\left\{p_{j}\right\}\right)$ if such a $C_{j}, p_{j}, C_{j^{\prime}}$ and $p_{j^{\prime}}$ exist. This is necessary and sufficient. We achieve this via a dynamic programming algorithm that fills a table $T$ where the cells are indexed by tuples of the form ( $C_{j}, p_{j}, p_{j^{\prime}}, r, Z$ ). Moreover, any non-empty cell indexed by the tuple $\left(C_{j}, p_{j}, p_{j^{\prime}}, r, Z\right)$ contains a single element of $\zeta\left(C_{j}, p_{j}, p_{j^{\prime}}, r, Z\right)$ and an empty cell indexed by the tuple ( $C_{j}, p_{j}, p_{j^{\prime}}, r, Z$ ) indicates that $\zeta\left(C_{j}, p_{j}, p_{j^{\prime}}, r, Z\right)$ is empty. Notice that if the table is filled correctly, then the solution (i.e., a NE strategy) can be determined by going over all possible $C_{j}, p_{j}, C_{j^{\prime}}$ and $p_{j^{\prime}}$ (of which there are polynomially many) and examining the entry of $T$ indexed by the tuple ( $C_{j}, p_{j}, p_{j^{\prime}}$, $\left.\log \ell,\left\{p_{j}\right\}\right)$.

We next describe how to fill the table $T$. We proceed by iteratively increasing the value of $r$ and in each iteration, filling all cells of $T$ corresponding to the value of $r$ in the current iteration. In the base case, $r=1$ and so for every $C_{j} \in C$, it follows that $\left|C_{r}\left(C_{j}\right)\right|=2$. Hence, for every $Z \in\binom{C}{\log \ell}, p_{j} \in C_{j}$ and $p_{j^{\prime}} \in C_{j^{\prime}} \in C_{1}\left(C_{j}\right) \backslash\left\{C_{j}\right\}$, it is straightforward to decide whether $\zeta\left(C_{j}, p_{j}, p_{j^{\prime}}, 1, Z\right) \neq \emptyset$ in polynomial time by simply trying all possible profiles. If it is non-empty, then we compute and add to $T\left(C_{j}, p_{j}, p_{j^{\prime}}, 1, Z\right)$ an arbitrary element of $\zeta\left(C_{j}, p_{j}, p_{j^{\prime}}, 1, Z\right)$. Otherwise we set $T\left(C_{j}, p_{j}, p_{j^{\prime}}, 1, Z\right)=\emptyset$ (including for those indices which $\zeta$ maps to $\emptyset$ by definition). Hence we may assume that we have filled the table $T$ for all entries with $r=1$. Note that this step takes time $n^{O(\log \ell)}$ since we have polynomially many choices for $C_{j}, p_{j}, C_{j^{\prime}}, p_{j^{\prime}}$ and $n^{O(\log \ell)}$ possibilities for $Z$ and furthermore, determining the entry $T\left(C_{j}, p_{j}, p_{j^{\prime}}, 1, Z\right)$ for each fixed choice of $C_{j}, p_{j}, C_{j^{\prime}}, p_{j^{\prime}}$ and $Z$ as described above takes only polynomial time.

Now, suppose that $r>1$ and inductively assume that for all $r^{\prime}<r$, for all choices of $C_{j}, C_{j^{\prime}}, p_{j}, p_{j^{\prime}}$, and $Z$, we have filled the table entry $T\left(C_{j}, p_{j}, p_{j^{\prime}}, r^{\prime}, Z\right)$ correctly. Now, fix a choice of $C_{j}, C_{j^{\prime}}$, $p_{j} \in C_{j}, p_{j^{\prime}} \in C_{j^{\prime}}$, and $Z \in\binom{\mathcal{P}}{\leq \log \ell}$ and we describe our procedure to fill the table entry $T\left(C_{j}, p_{j}, p_{j^{\prime}}, r, Z\right)$. We check if there is a $\hat{s}_{1} \in T\left(C_{j}, p_{j}, p_{j_{1}}, r-1, Z \cup\left\{p_{j^{\prime}}\right\}\right)$ and a $\hat{s}_{2} \in T\left(C_{j^{\prime}}, p_{j^{\prime}}, p_{j_{2}}, r-\right.$ $\left.1, Z \cup\left\{p_{j^{\prime}}\right\}\right)$ for some choice of $p_{j_{1}}$ and $p_{j_{2}}$. If yes, then we set $T\left(C_{j}, p_{j}, p_{j^{\prime}}, r, Z\right)=\hat{s}_{1} \cdot \hat{s}_{2}$ and otherwise we set it to $\emptyset$. The second point of Lemma 2 indicates that composing $\hat{s}_{1}$ and $\hat{s}_{2}$ in this way indeed results in a profile that is contained in $\zeta\left(C_{j}, p_{j}, p_{j^{\prime}}, r, Z\right)$, implying the correctness of our algorithm. Finally, the claimed running time bound follows from the fact that the table $T$ has $n^{O(\log \ell)}$ entries in total, each of which is being filled in polynomial time using constant-time lookups into polynomially-many previously filled entries of $T$.

Observe that our algorithm actually achieves more than what is claimed in the statement of the theorem. Indeed, for every player $p_{j}$, our algorithm computes an equilibrium in which $p_{j}$ is the winner (if one exists).

### 3.2 Beyond Win-Lose Games

In earlier sections we assumed coalitions to be only interested in winning the tournament. It is however natural to look at scenarios in which participants have a major incentive to progress as far as possible even if they do not win, as in the UEFA Champions League.

Here, we study coalitional games in which the teams' utility is determined by the round they reach. We call these games beyond win/lose (B-W/L) games. Formally, we define utility functions $u_{1}, \ldots, u_{\ell}$ where for any $s=\left(p_{1}, \ldots, p_{\ell}\right), u_{i}(s)=k$ if and only $p_{i} \in S E_{s}^{k} \backslash S E_{s}^{k+1}$. As with the win/lose scenario, in a NE no coalition can improve their utility by selecting different player.

Definition 5 (Beyond Win/Lose NE). A profile $s=\left(p_{1}, \ldots, p_{\ell}\right)$ is $a$ : Beyond Win-Lose Nash Equilibrium ( $B-W / L N E$ ), if for all $i, k \in[\log \ell]_{\geq 0}, p_{i}^{\prime}:$ if $u_{i}\left(s\left[p_{i}^{\prime}\right]\right)=k$, then $u_{i}(s) \geq k$.

In the remainder of this section we focus on structural results characterising B-W/L NE and algorithmic results addressing the problems of recognition and computation for each.
3.2.1 Algorithmic properties of $B-W / L N E$. We begin by providing a structural result which we then use for recognising and computing B-W/L NEs, capturing when a profile is a B-W/L NE.

Lemma 3. Consider a profile $s=\left(p_{i}, \ldots, p_{\ell}\right)$ and let $j \in[\ell]$ such that $p_{j}$ wins $S E_{s}$. Furthermore, suppose that $p_{j^{\prime}} \in C_{j^{\prime}}$ is the final opponent of $p_{j}$ in $S E_{s}$. Then, $s$ is a $B-W / L N E$ if and only if (a) $s_{1}=\left.s\right|_{C_{\log \ell-1}\left(C_{j}\right)}$ and $s_{2}=\left.s\right|_{C \backslash C_{\log \ell-1}\left(C_{j}\right)}$ are both B-W/L NEs and (b) there is no $p_{j^{\prime}}^{\prime} \in C_{j^{\prime}}$ which beats $p_{j}$ such that $p_{j^{\prime}}^{\prime}$ wins $S E_{s_{2}\left[p_{j^{\prime}}^{\prime}\right]}^{\log \ell-1}$.

Proof. In the forward direction, suppose that $s$ is a B-W/L NE. Observe that if Property (b) is violated then $C_{j^{\prime}}$ can improve its position in $S E_{s\left[p_{j}^{\prime}\right]}^{\prime}$, so $s$ is not a B-W/L NE. Moreover, suppose that one of $s_{1}$ or $s_{2}$ is not a B-W/L NE for the respective subtournaments. W.l.o.g., suppose that for some $C_{i} \in C \backslash\left\{C_{j}, C_{j^{\prime}}\right\}, u_{i}\left(s_{1}\left[p_{i}^{\prime}\right]\right)>u_{i}\left(s_{1}\right)$. That is, $C_{i}$ is able to improve its position in the subtournament played by $C_{\log \ell-1}\left(C_{j}\right)$ by playing $p_{i}^{\prime}$ instead of $p_{i}$. Then, $u_{i}\left(s\left[p_{i}^{\prime}\right]\right)>$ $u_{i}(s)$. That is, $C_{i}$ can also strictly improve its position in $s\left[p_{i}^{\prime}\right]$, a contradiction to $s$ being a B-W/L NE. The argument when $C_{i} \in$ $C \backslash C_{\log \ell-1}\left(C_{j}\right)$ is identical and hence Property (a) also holds.

Conversely, suppose that Properties (a) and (b) hold and $s$ is not a B-W/L NE. Let $p_{i}^{\prime} \in C_{i}$ be such that $u_{i}\left(s\left[p_{i}^{\prime}\right]\right)>u_{i}(s)$. Since Property (b) holds, it cannot be the case that $i \in\left\{j, j^{\prime}\right\}$. Moreover, if $C_{i} \in C^{\log \ell-1}\left(C_{j}\right)$, then it contradicts our assumption that $s_{1}$ is a B-W/L NE and, otherwise, our assumption that $s_{2}$ is one.

This allows us to provide a P-time algorithm for recognising a B-W/L NE, which can be extended to the computation problem.

## Proposition 3. Recognising a $B-W / L$ NE is $P$-time solvable.

Proof. By Lemma 3, we have that for a given profile $s=\left(p_{i}, \ldots, p_{\ell}\right)$ and $j \in[\ell]$ such that $p_{j}$ is the winner of $S E_{s}, s$ is a B-W/L NE if and only if there is no $j^{\prime} \in[\ell] \backslash\{j\}$ and player $p_{j^{\prime}}^{\prime} \in C_{j^{\prime}}$ such that $p_{j^{\prime}}^{\prime}$ beats the parent of $p_{j^{\prime}}$ as well as all players in the set Opponent ${ }^{+}\left[p_{j^{\prime}}, p_{j^{\prime}}\right]$. The P-time algorithm for recognising a B-W/L NE follows.

Theorem 2. Computing a B-W/L NE is P-time solvable.

Proof. Notice that if there is a B-W/L NE, then for $r=1$, for every $C_{j} \in C$, the subtournament over $C_{r}\left(C_{j}\right)$ can be won by exactly one out of the two coalitions in this subtournament and the set of potentially winning players that can participate in a B-W/L NE can be easily computed in polynomial time. Indeed, suppose that $C_{j}$ and $C_{j^{\prime}}$ are the two coalitions in this subtournament. If every player in $C_{j}$ loses to some player in $C_{j^{\prime}}$ and every player in $C_{j^{\prime}}$ loses to some player in $C_{j}$, then in any profile, at least one out of these two coalitions will be able to improve their final position by at least one place, implying the non-existence of a B-W/L NE. Hence, we identify all players in $C_{j}\left(C_{j^{\prime}}\right)$ that beats every player in $C_{j^{\prime}}$ (respectively, $C_{j}$ ) by examining all matches between players in these two coalitions. This is clearly P-time computable.

We now inductively argue a similar property for $r>1$. By the induction hypothesis, we have that if there is a B-W/L NE, then for every $C_{j} \in C$, the subtournaments over $C_{r}\left(C_{j}\right) \backslash C_{r-1}\left(C_{j}\right)$ and $C_{r-1}\left(C_{j}\right)$ can be won by exactly one coalition each (say, $C^{\prime}$ and $C^{\prime \prime}$ respectively) and the potentially winning players from each of these coalitions (denoted by $\mathcal{P}^{\star}\left(C^{\prime}\right)$ and $\mathcal{P}^{\star}\left(C^{\prime \prime}\right)$ respectively) that can participate in a B-W/L NE can be computed in polynomial time. Now, notice that for a B-W/L NE to exist, exactly one out of the following two cases must occur: either (i) there is a player in $\mathcal{P}^{\star}\left(C^{\prime}\right)$ that beats every player in $\mathcal{P}^{\star}\left(C^{\prime \prime}\right)$ or (ii) there is a player in $\mathcal{P}^{\star}\left(C^{\prime \prime}\right)$ that beats every player in $\mathcal{P}^{\star}\left(C^{\prime}\right)$. Moreover, the set of players satisfying (i) or (ii) can be computed in polynomial time. $\quad \square$

## 4 DYNAMIC KNOCKOUT TOURNAMENTS

We now allow coalitions to choose players at each round of the tournament. In this model, a strategy of a coalition $C$ consists of $\ell-1$ (not necessarily distinct) players representing, for each opposing coalition, a choice of a player of $C$ to face said opposing coalition. We therefore model coalitional choices as dynamic strategy profiles $\sigma: C \rightarrow(C \rightarrow \mathcal{P})$. Specifically, for each distinct $i, j \in[\ell], \sigma\left(C_{i}\right)\left(C_{j}\right)$ elects a member of the coalition $C_{i}$ when facing coalition $C_{j}$. For every $C \in \mathcal{C}$, let $\mathcal{S}(C)$ denote the set of all functions $C \rightarrow C$, i.e. the set of all possible responses to a given coalition. We require dynamic strategy profiles $\sigma$ to be such that $\sigma\left(C_{i}\right) \in \mathcal{S}\left(C_{i}\right)$ for each $C_{i} \in C$. We use $\rho_{i}$ to denote $\sigma\left(C_{i}\right)$ for every coalition $C_{i} \in C$. Equivalently, we represent $\sigma$ as a tuple $\left(\rho_{1}, \ldots, \rho_{\ell}\right)$ where $\rho_{i} \in$ $\mathcal{S}\left(C_{i}\right)$, and denote with $\rho_{i j}$ the player in $C_{i}$ selected to play against $C_{j}$ (also named $\rho_{i}\left(C_{j}\right)$ ). In this interpretation, for every $i \in[\ell]$, $\sigma\left(C_{i}\right)\left(C_{i}\right)$ is meaningless as a coalition does not face itself, and so we assigned an arbitrary player of $C_{i}$ to be $\sigma\left(C_{i}\right)\left(C_{i}\right)$ (say the lexicographically smallest one).

Input representation. Notice that a dynamic strategy can be represented as a matrix $M$ of the size $\ell \times \ell$, where $\ell$ is the number of coalitions. Then, each entry of the matrix $M[i, j]$ such that $i \neq j$ corresponds to $\sigma\left(C_{i}\right)\left(C_{j}\right)$. If $i=j, M[i, j]$ is empty. Given this representation, a dynamic strategy can be encoded in space $O\left(\ell^{2} \log n\right)$.

While dynamic games typically warrant history-dependent strategies, the tree structure of coalitional knockout tournaments allows for a pair of coalitions to meet only at a unique round, if ever, as a function of the initial seeding. Therefore, each coalitional choice is effectively history-dependent when the initial seeding is fixed.

Mirroring the one-shot case, we represent tournaments between coalitions, as induced from a dynamic strategy profile. Let $D_{C}$ be a
digraph on $\mathcal{P}$ and $\sigma$ a dynamic strategy profile over $\mathcal{C}$. A dynamic coalitional digraph $\left(D_{C}, \sigma\right)$ is a digraph with vertex set $C$ where $(i, j)$ is an arc of $\left(D_{C}, \sigma\right)$ if and only if $\left(\sigma_{i}\left(C_{j}\right), \sigma_{j}\left(C_{i}\right)\right)$ is an arc of $D_{C}$. Also, the notion of binomial arborescence $S E_{\sigma}$ is lifted from the one-shot case in the natural manner.

Further, the solution concept is lifted from the one-shot case as expected, both for the win-lose ( $\mathrm{d}-\mathrm{NE}$ ) and the beyond win-lose (BWL d-NE) games. The techniques used in the proofs of the results in this section are similar to their one-shot correspondents. Thus, we omit some of them.

### 4.1 Dynamic Win-Lose Games

In this section, we begin by lifting NEs to the current setting.
Definition 6. A dynamic profile $\sigma=\left(\rho_{1}, \ldots, \rho_{\ell}\right)$ is a dynamic Nash Equilibrium ( $d-N E$ ), if for all i and for all $\rho_{i}^{\prime}:$ if $S E_{\pi, \sigma\left[\rho_{i}^{\prime}\right]}^{\log \ell}=C_{i}$ then $S E_{\sigma}^{\log \ell}=C_{i}$.

So, a dynamic strategy profile is a dynamic NE if no losing coalition can become a winner by changing their strategy (which now corresponds to selecting a player for each opposing coalition) unilaterally. We establish a useful structural equivalence between d-NE, and NE for a one-shot coalitional tournament defined on an auxiliary tournament digraph where the players correspond to the set of the given parties' strategies. We first define the auxiliary digraph.

Definition 7. Consider the digraph $D_{C}$. We define by $D_{C}^{D y n}$ the graph with vertex set $\bigcup_{C \in C} \mathcal{S}(C)$ and arc set defined as follows. For every distinct $C_{i}, C_{j} \in \mathcal{C}$, for every $\rho_{i} \in \mathcal{S}\left(C_{i}\right)$ and $\rho_{j} \in \mathcal{S}\left(C_{j}\right)$, there is an $\operatorname{arc}\left(\rho_{i}, \rho_{j}\right)$ if $\left(\rho_{i}\left(C_{j}\right), \rho_{j}\left(C_{i}\right)\right)$ is an arc in $D_{C}$ and there is an $\operatorname{arc}\left(\rho_{j}, \rho_{i}\right)$ otherwise.

That is, $D_{C}^{\text {Dyn }}$ has the arc $\left(\rho_{i}, \rho_{j}\right)$ if and only if the player $p_{i}$ elected to face $C_{j}$ by $C_{i}$ wins against the player $C_{j}$ elected to face $C_{i}$. Observe that like $D_{C}, D_{C}^{\text {Dyn }}$ is also a $|\mathcal{C}|$-partite tournament digraph with a partition for every $\mathcal{S}\left(C_{i}\right)$ where $C_{i} \in C$.

This construction allows us to reason about the dynamic solution concepts in terms of one-shot tournaments, as the following characterisation shows.

Lemma 4. Let $\sigma$ be a dynamic strategy profile over the coalitions $C$. Then, $\sigma$ is a d-NE for the dynamic SE-tournament over $C$ using the pairwise results in $D_{C}$ and seeding $\pi$ iff $\left(\rho_{1}, \ldots, \rho_{\ell}\right)$ is a NE for the one-shot SE-tournament over the coalitions $\mathcal{S}(C)=\{\mathcal{S}(C) \mid C \in C\}$ using the pairwise results in $D_{C}^{D y n}$ and seeding $\pi$ projected from the coalitions in $C$ to those in $\mathcal{S}(C)$ in the natural way.
4.1.1 Algorithmic properties of $d-N E$. We now address the algorithmic questions of recognising and computing a dynamic NE. Observe that in $D_{C}^{\mathrm{Dyn}}$, the number of vertices could be as large as $\ell \cdot m^{\ell}$ where $m$ is the size of the largest coalition in $D_{C}$. As a result, although we can transfer our structural results on NEs from the one-shot setting to the dynamic setting using the graph $D_{C}^{\text {Dyn }}$, we cannot simply use the same algorithms because the running time, though polynomial in the size of $D_{C}^{\mathrm{Dyn}}$, will no longer remain polynomial in the size of the actual input, which is linear in the size of the graph $D_{C}$. However, using appropriate queries that can
be answered in polynomial-time (in the size of $D_{C}$ ), we can still obtain a polynomial-time algorithm for recognising a d-NE. In the rest of the paper, polynomial-time refers to polynomial-time in the size of $D_{C}$.

Lemmata 1 and 4 entail the tractability of recognising d-NE.
Proposition 4. Recognising a d-NE is P-time solvable.
Proof. By invoking Lemma 1 on $D_{C}^{\text {Dyn }}$ and the equivalence given by Lemma 4, we conclude that a given dynamic profile $\sigma=$ ( $\rho_{1}, \ldots, \rho_{\ell}$ ) won by $C_{j}$ is a d-NE if and only if there is no $j^{\prime} \in$ $[\ell] \backslash\{j\}$ and strategy $\rho_{j^{\prime}}^{\prime} \in \mathcal{S}\left(C_{j^{\prime}}\right)$ such that $\rho_{j^{\prime}}^{\prime}$ beats all players in the set Opponent ${ }^{+}\left[\rho, \rho_{j^{\prime}}\right]$ for every player $\rho$ on the $\rho_{j^{-}} \rho_{j^{\prime}}$ path in the b.a $S E_{\sigma}$ contained in the graph $D_{C}^{\mathrm{Dyn}}$. Observe that given $\sigma$, the b.a. $\mathrm{d}-\mathrm{SE} E_{\sigma}$ and the set Opponent ${ }^{+}\left[\rho, \rho_{j^{\prime}}\right]$ for every player $\rho$ on the $\rho_{j}-\rho_{j^{\prime}}$ path can be computed in polynomial time by simply querying, for every pair of coalitions - "who is the winner between the two" $i f$ both played their respective strategies contained in $\sigma$. Now, for any $j^{\prime} \in[\ell] \backslash\{j\}$, we can check whether there is a strategy $\rho_{j^{\prime}}^{\prime} \in \mathcal{S}\left(C_{j^{\prime}}\right)$ which beats the players selected by the respective coalitions by simply inspecting the arcs of the graph $D_{C}$.

Then, with reasoning similar to the one-shot case we obtain a quasi-polynomial-time algorithm for computing a $d$-NE if it exists.

Theorem 3. A $d$-NE can be computed in $n^{O(\log \ell)}$-time.

### 4.2 Beyond Dynamic Win-Lose Games

Let us further analyse the final case, that of tournaments where coalitions can modify their choices at each round and are interested in tournament progression. The solution studied in this section is a natural modification of that considered earlier.

We define utility functions $u_{1}, \ldots, u_{\ell}$ where, for every profile $\sigma=\left(\rho_{1}, \ldots, \rho_{\ell}\right), u_{i}(\sigma)=k$ if and only if $C_{i} \in S E_{\sigma}^{k} \backslash S E_{\sigma}^{k+1}$. That is, Coalition $C_{i}$ wins $k$ rounds but not $k+1$ rounds. This means that the utility of coalition $i$ is dependent on the final round which they reach under a given strategy profile. The solution concept mirrors the one we defined previously.)

Definition 8. A dynamic profile $\sigma=\left(\rho_{1}, \ldots, \rho_{\ell}\right)$ is $a$ Beyond Win-Lose Dynamic Nash Equilibrium ( $B-W / L d-N E$ ), if for all $i, k \in$ $[\log \ell]_{\geq 0}, \rho_{i}^{\prime}:$ if $u_{i}\left(\sigma\left[\rho_{i}^{\prime}\right]\right)=k$, then $u_{i}(\sigma) \geq k$.
4.2.1 Algorithmic properties of $B-W / L d-N E$. Let us commence by the study of NE in the beyond win/lose, dynamic setting. We first provide a structural result on which we base the analysis of its algorithmic properties.

Lemma 5. Consider a dynamic profile $\sigma=\left(\rho_{i}, \ldots, \rho_{\ell}\right)$ and let $j \in[\ell]$ such that $C_{j}$ wins $S E_{\sigma}$. Furthermore, suppose that $C_{j^{\prime}}$ is the final opponent of $C_{j}$ in $S E_{\sigma}$. Then, $\sigma$ is a $B-W / L d-N E$ if and only if $\sigma_{1}=\left.\sigma\right|_{C^{\log \ell-1}\left(C_{j}\right)}$ and $\sigma_{2}=\left.\sigma\right|_{C \backslash C^{\log \ell-1}\left(C_{j}\right)}$ are both B-W/L d-NEs and $(b)\left(\rho_{j}, \rho_{j^{\prime}}\right)$ is an $N E$.

Proof. In the forward direction, consider a B-W/L d-NE $\sigma=$ $\left(\rho_{i}, \ldots, \rho_{\ell}\right)$. Observe that if Property (b) is violated then $C_{j^{\prime}}$ can improve its position by choosing $\rho_{j^{\prime} j} \in C_{j^{\prime}}$ as a player which beats $\rho_{j j^{\prime}}$, contradicting our assumption that $\sigma$ is a B-W/L d-NE. Hence we conclude that Property (b) is satisfied. On the other

ONE-SHOT DYNAMIC

|  | CHECK | FIND | CHECK | FIND |
| :--- | :---: | :---: | :---: | :---: |
| W/L NE | P | $\mathrm{Q}-\mathrm{P}$ | P | $\mathrm{Q}-\mathrm{P}$ |
| B W/L NE | P | P | P | P |

Table 1: Summary of the algorithmic results in the paper.
hand, suppose that one of $\sigma_{1}$ or $\sigma_{2}$ is not a B-W/L d-NE for the respective subtournaments. Without loss of generality, suppose that for some $C_{i} \in \mathcal{C} \backslash\left\{C_{j}, C_{j^{\prime}}\right\}, C_{i}$ is able to improve its position in the subtournament played by $C_{\log \ell-1}\left(C_{j}\right)$ by choosing the strategy $\rho_{i}^{\prime}$ instead of $\rho_{i}$. Then, $C_{i}$ can also strictly improve its position in $S E_{\sigma\left[\rho_{i}^{\prime}\right]}$, a contradiction to $\sigma$ being a B-W/L d-NE. The case when $C_{i} \in C \backslash C_{\log \ell-1}\left(C_{j}\right)$ is identical and hence Property (a) holds.

Conversely, suppose that Properties (a) and (b) hold and suppose that $\sigma$ is not a B-W/L NE. Let $\rho_{i}^{\prime} \in \mathcal{S}\left(C_{i}\right)$ be such that $u_{i}\left(\sigma\left[\rho_{i}^{\prime}\right]\right)>$ $u_{i}(\sigma)$. Since Property (b) holds, it cannot be the case that $i \in\left\{j, j^{\prime}\right\}$. Moreover, if $C_{i} \in C^{\log \ell-1}\left(C_{j}\right)$, then it contradicts our assumption that $\sigma_{1}$ is a B-W/L d-NE and otherwise, it contradicts our assumption that $\sigma_{2}$ is a B-W/L d-NE.

The following algorithmic result follows from Lemma 5.
Theorem 4. Recognising and computing a $B$-W/L $d$-NE are $P$-time solvable.

Notice that a coalitional knockout tournament can admit multiple dynamic B-W/L NEs for a fixed seeding, but these are all outcome-equivalent.

## 5 DISCUSSION

We introduced a model for knockout tournaments played between groups, to allow for each group to strategically select a member to take part in it. We carried out an algorithmic analysis of Nash equilibrium strategies under various setups occurring in practice, showing tractable results (polynomial-time or quasi-polynomialtime) for all cases. Table 1 gives an overview of our results.

We covered, for space reasons, only the treatment of Nash equilibria, but alternative solution concepts have also been studied. In particular, dominant strategy equilibria (DSE), where a profile $s=\left(p_{1}, \ldots, p_{\ell}\right)$ is called a DSE if for all $i$ and $s^{\prime}:$ if $C_{i}$ wins $S E_{s^{\prime}} \in C_{i}$ then $C_{i}$ wins $S E_{s^{\prime}\left[p_{i}\right]}$. In other words, no coalition can improve on their DSE strategy, irrespective of the choices of their opponents. Adapting this one-shot win-lose definition to the beyond win-lose and dynamic counterparts, analogously to what done with Nash equilibria, P-time algorithms can be shown for all cases in Table 1.

We foresee various potential research directions building on our work, which involve relaxing some of the assumptions, especially on participants knowledge and strategies, and looking at alternative solution concepts. Here we discuss a few concrete ones.
(1) It is interesting to explore fast parameterized algorithms for finding a one-shot or dynamic W/L NE, e.g., parameterized by the number of coalitions $\ell$.
(2) All our results so far have assumed a tournament with a fixed seeding. However the choice of seeding may influence the tournament enormously, in particular some seedings may admit an equilibrium while others may not. Establishing the complexity of finding a seeding such that a given solution concept exists is therefore an important problem in this setup. This has repercussions for tournament fixing [26], as a malicious external attacker may be in a position to choose a seeding with a favourable winner in all resulting equilibria and high complexity barriers may act against it.
(3) A third direction for future research is to establish the existential and algorithmic properties of the considered solution concepts under relaxed versions of our model. For instance, consider the scenario in which the beating relation or tournament digraph is relaxed to be stochastic. In this case, studying the existence and computational complexity of various equilibria based on the expected utility would be of high interest. This would bring us closer to understanding a setting that models real-world scenarios more faithfully than possible using only a "static" tournament digraph.
(4) The combination of sequential decision-making in tournaments and the beating relation between players suggests novel solution concepts. Assume that coalition $A$ can win a win-lose one-shot tournament, provided coalition $B$ does not field player $b$, who defeats every player in $A$. We have that, every time $B$ fields $b, A$ is indifferent to the choice of any of their players. However there is a sense in which $A$ must play a potential winner, should something happen to $b$. This suggests a trembling-hand interpretation of the beating relation (building on a classical solution concept [25]), which coalitions can try and exploit. Intuitively, profile $s$ is a trembling-hand tournament perfect equilibrium if, for every coalition $C_{i}$, the strategy profile $s$ is a Nash equilibrium at each sub-tournament, where $C_{i}$ plays $s_{i}$ as if they were able to reach that sub-tournament.
(5) Another natural research direction concerns the possibility of cooperative and semi-cooperative behaviour among coalitions. Assume that coalition $A$ needs coalition $B$ to win the tournament, but $B$ can never win. We envisage an interesting variant of endogenous games [15] arising in our tournaments. Before the tournament starts, $A$ can transfer a part of the expected payoff to $B$, should $B$ refrain from fielding a player blocking $A$. This should be seen as a form of manipulation, e.g., flipping results of a certain number of matches, where incentives comes from the players themselves
(6) Finally, the extension to mixed strategies is a natural followup. It is then interesting to explore whether the knockout tournament structure can add any advantage in terms of (mixed-strategy) Nash equilibrium computation, as opposed to full-blown normal form games.

## ACKNOWLEDGEMENTS

This work was supported by the Engineering and Physical Sciences Research Council [grant number EP/V007793/1].

## REFERENCES

[1] Alon Altman, Ariel D Procaccia, and Moshe Tennenholtz. 2009. Nonmanipulable selections from a tournament. In Twenty-First International foint Conference on Artificial Intelligence. IJCAI, 686-690.
[2] Haris Aziz, Serge Gaspers, Simon Mackenzie, Nicholas Mattei, Paul Stursberg and Toby Walsh. 2014. Fixing a balanced knockout tournament. In Twenty-Eighth AAAI Conference on Artificial Intelligence. AAAI, 552-558.
[3] Haris Aziz, Serge Gaspers, Simon Mackenzie, Nicholas Mattei, Paul Stursberg, and Toby Walsh. 2018. Fixing balanced knockout and double elimination tournaments. Artificial Intelligence 262 (2018), 1-14.
[4] Felix Brandt, Markus Brill, and Bernhard Harrenstein. 2016. Tournament Solutions. In Handbook of Computational Social Choice, F Brandt, V Conitzer, U Endriss, J Lang, and A. D. Procaccia (Eds.). Cambridge University Press, 453-474
[5] Felix Brandt, Markus Brill, and Paul Harrenstein. 2018. Extending tournament solutions. Social Choice and Welfare 51, 2 (2018), 193-222.
[6] Felix Brandt, Markus Brill, Hans Georg Seedig, and Warut Suksompong. 2018. On the structure of stable tournament solutions. Economic Theory 65, 2 (2018) 483-507.
[7] Felix Brandt and Felix A. Fischer. 2008. Computing the minimal covering set. Mathematical Social Sciences 56, 2 (2008), 254-268
[8] Felix Brandt, Felix A. Fischer, Paul Harrenstein, and Maximilian Mair. 2010. A computational analysis of the tournament equilibrium set. Social Choice and Welfare 34, 4 (2010), 597-609.
[9] Markus Brill and Vincent Conitzer. 2015. Strategic Voting and Strategic Candidacy. In Twenty-Ninth AAAI Conference on Artificial Intelligence. AAAI, 819-826.
[10] Krishnendu Chatterjee, Rasmus Ibsen-Jensen, and Josef Tkadlec. 2016. Robust Draws in Balanced Knockout Tournaments. In Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence. IJCAI, 172-179.
[11] Piotr Faliszewski, Laurent Gourvès, Jérôme Lang, Julien Lesca, and Jérôme Monnot. 2016. How hard is it for a party to nominate an election winner?. In TwentyFifth International foint Conference on Artificial Intelligence. IJCAI, 257-263.
[12] David Fisher and Jennifer Ryan. 1995. Tournament games and positive tournaments. Journal of Graph Theory 19, 2 (1995), 217-236.
[13] Paul Harrenstein, Grzegorz Lisowski, Ramanujan Sridharan, and Paolo Turrini. 2021. A Hotelling-Downs Framework for Party Nominees. In Twentieth International Conference on Autonomous Agents and Multiagent Systems. AAMAS,

593-601.
[14] Noam Hazon, Paul Dunne, Sarit Kraus, and Michael Wooldridge. 2008. How to rig elections and competitions. In Second International Workshop on Computational Social Choice. COMSOC, 301-312.
[15] Matthew O. Jackson and Simon Wilkie. 2005. Endogenous Games and Mechanisms: Side Payments Among Players. Review of Economic Studies 72, 2 (2005), 543-566.
[16] Michael P Kim, Warut Suksompong, and Virginia Vassilevska Williams. 2017. Who can win a single-elimination tournament? SIAM Journal on Discrete Mathematics 31, 3 (2017), 1751-1764.
[17] Aleksei Y. Kondratev and Vladimir V. Mazalov. 2019. Tournament solutions based on cooperative game theory. International fournal of Game Theory 49, 1 (2019), 119-145.
[18] Christine Konicki and Virginia Vassilevska Williams. 2019. Bribery in Balanced Knockout Tournaments. In Eighteenth International Conference on Autonomous Agents and Multi Agent Systems. AAMAS, 2066-2068.
[19] Gilbert Laffond, Jean Lainé, and Jean-François Laslier. 1996. Compositionconsistent tournament solutions and social choice functions. Social Choice and Welfare 13, 1 (1996), 75-93.
[20] G. Laffond, J. F. Laslier, and M. Le Breton. 1993. The bipartisan set of a tournament game. Games and Economic Behavior 5 (1993), 182-201. Issue 1.
[21] Jean-Francois Laslier. 1997. Tournament solutions and majority voting. Springer Verlag.
[22] Reshef Meir. 2018. Strategic Voting. Morgan \& Claypool.
[23] JW Moon. 1968. Topics on tournaments. Holt, Reinhart and Winston.
[24] Hervé Moulin. 1986. Choosing from a tournament. Social Choice and Welfare 3, 4 (1986), 271-291.
[25] Reinhard Selten. 1975. Reexamination of the perfectness concept for equilibrium points in extensive games. International fournal of Game Theory 4 (1975), 25-55.
[26] Virginia Vassilevska Williams. 2010. Fixing a Tournament. In Twenty-Fourth AAAI Conference on Artificial Intelligence. AAAI, 895-900.
[27] Thuc Vu, Alon Altman, and Yoav Shoham. 2009. On the complexity of schedule control problems for knockout tournaments. In Eigth International Conference on Autonomous Agents and Multiagent Systems. AAMAS, 225-232.


[^0]:    Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), P. Faliszewski, V. Mascardi, C. Pelachaud, M.E. Taylor (eds.), May 9-13, 2022, Online. © 2022 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

