

Rankings for Bipartite Tournaments via Chain Editing

Joseph Singleton
Cardiff University
Cardiff

singletonj1@cardiff.ac.uk

Richard Booth
Cardiff University
Cardiff

boothr2@cardiff.ac.uk

ABSTRACT

Ranking the participants of a tournament has applications in voting, paired comparisons analysis, sports and other domains. In this paper we introduce *bipartite tournaments*, which model situations in which two different kinds of entity compete indirectly via matches against players of the opposite kind; examples include education (students/exam questions) and solo sports (golfers/courses). In particular, we look to find rankings via *chain graphs*, which correspond to bipartite tournaments in which the sets of adversaries defeated by the players on one side are nested with respect to set inclusion. Tournaments of this form have a natural and appealing ranking associated with them. We apply *chain editing* – finding the minimum number of edge changes required to form a chain graph – as a new mechanism for tournament ranking. The properties of these rankings are investigated in a probabilistic setting, where they arise as maximum likelihood estimators, and through the axiomatic method of social choice theory. Despite some nice properties, two problems remain: an important anonymity axiom is violated, and chain editing is NP-hard. We address both issues by relaxing the minimisation constraint in chain editing, and characterise the resulting ranking methods via a greedy approximation algorithm.

KEYWORDS

Social Choice; Tournaments; Ranking; Chain Graphs

ACM Reference Format:

Joseph Singleton and Richard Booth. 2021. Rankings for Bipartite Tournaments via Chain Editing. In *Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), Online, May 3–7, 2021*, IFAAMAS, 9 pages.

1 INTRODUCTION

A tournament consists of a finite set of players equipped with a *beating relation* describing pairwise comparisons between each pair of players. Determining a ranking of the players in a tournament has applications in voting in social choice [5] (where players represent alternatives and x beats y if a majority of voters prefer x over y), paired comparisons analysis [14] (where players may represent products and the beating relation the preferences of a user), search engines [24], sports tournaments [3] and other domains.

In this paper we introduce *bipartite tournaments*, which consist of two disjoint sets of players A and B such that comparisons only take place between players from opposite sets. We consider ranking methods which produce two rankings for each tournament – one for each side of the bipartition. Such tournaments model situations in which two different kinds of entity compete *indirectly* via matches

against entities of the opposite kind. The notion of competition may be abstract, which allows the model to be applied in a variety of settings. An important example is education [15], where A represents students, B exam questions, and student a ‘beats’ question b by answering it correctly. Here the ranking of students reflects their performance in the exam, and the ranking of questions reflects their *difficulty*. The simultaneous ranking of both sides allows one ranking to influence the other; e.g. so that students are rewarded for correctly answering difficult questions. This may prove particularly useful in the context of crowdsourced questions provided by students themselves, which may vary in their difficulty (see for example the PeerWise system [8]).

A related example is *truth discovery* [16, 21]: the task of finding true information on a number of topics when faced with conflicting reports from sources of varying (but unknown) reliability. Many truth discovery algorithms operate iteratively, alternately estimating the reliability of sources based on current estimates of the true information, and obtaining new estimates of the truth based on source reliability levels. The former is an instance of a bipartite tournament; similar to the education example, A represents data sources, B topics of interest, and a defeats b by providing true information on topic b (according to the current estimates of the truth). Applying a bipartite tournament ranking method at this step may therefore facilitate development of *difficulty-aware* truth discovery algorithms, which reward sources for providing accurate information on difficult topics [12]. Other application domains include the evaluation of generative models in machine learning [19] (where A represents generators and B discriminators) and solo sports contests (e.g. where A represents golfers and B golf courses).

In principle, bipartite tournaments are a special case of *generalised* tournaments [7, 14, 23], which allow intensities of victories and losses beyond a binary win or loss (thus permitting draws or multiple comparisons), and drop the requirement that every player is compared to all others. However, many existing ranking methods in the literature do not apply to bipartite tournaments due to the violation of an *irreducibility* requirement, which requires that the tournament graph be strongly connected. In any case, bipartite tournament ranking presents a unique problem – since we aim to rank players with only indirect information available – which we believe is worthy of study in its own right.

In this work we focus particularly on ranking via *chain graphs* and *chain editing*. A chain graph is a bipartite graph in which the neighbourhoods of vertices on one side form a chain with respect to set inclusion. A (bipartite) tournament of this form represents an ‘ideal’ situation in which the capabilities of the players are perfectly nested: weaker players defeat a subset of the opponents that stronger players defeat. In this case a natural ranking can be formed according to the set of opponents defeated by each player. These rankings respect the tournament results in an intuitive sense:

Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), U. Endriss, A. Nowé, F. Dignum, A. Lomuscio (eds.), May 3–7, 2021, Online. © 2021 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

if a player a defeats b and b' ranks worse than b , then a must defeat b' also. Unfortunately, this perfect nesting may not hold in reality: a weak player may win a difficult match by coincidence, and a strong player may lose a match by accident. With this in mind, Jiao et al. [15] suggested an appealing ranking method for bipartite tournaments: apply *chain editing* to the input tournament – i.e. find the minimum number of edge changes required to form a chain graph – and output the corresponding rankings. Whilst their work focused on algorithms for chain editing and its variants, we look to study the properties of the ranking method itself through the lens of computational social choice.

Contribution. Our primary contribution is the introduction of a class of ranking mechanisms for bipartite tournaments defined by chain editing. We also provide a new probabilistic characterisation of chain editing via maximum likelihood estimation. To our knowledge this is the first in-depth study of chain editing as a ranking mechanism. Secondly, we introduce a new class of ‘chain-definable’ mechanisms by relaxing the minimisation constraint of chain editing in order to obtain tractable algorithms and to resolve the failure of an important anonymity axiom.

Paper outline. In Section 2 we define the framework for bipartite tournaments and introduce chain graphs. Section 3 outlines how one may use chain editing to rank a tournament, and characterises the resulting mechanisms in a probabilistic setting. Axiomatic properties are considered in Section 4. Section 5 defines a concrete scheme for producing chain-editing-based rankings. Section 6 introduces new ranking methods by relaxing the chain editing requirement. Related work is discussed in Section 7, and we conclude in Section 8. Note that some proofs are omitted due to lack of space, and can be found in the appendix of [22].

2 PRELIMINARIES

In this section we define our framework for bipartite tournaments, introduce chain graphs and discuss the link between them.

2.1 Bipartite Tournaments

Following the literature on generalised tournaments [7, 14, 23], we represent a tournament as a matrix, whose entries represent the results of matches between participants. In what follows, $[n]$ denotes the set $\{1, \dots, n\}$ whenever $n \in \mathbb{N}$.

Definition 2.1. A bipartite tournament – hereafter simply a *tournament* – is a triple (A, B, K) , where $A = [m]$ and $B = [n]$ for some $m, n \in \mathbb{N}$, and K is an $m \times n$ matrix with $K_{ab} \in \{0, 1\}$ for all $(a, b) \in A \times B$. The set of all tournaments will be denoted by \mathcal{K} .

Here A and B represent the two sets of players in the tournament.¹ An entry K_{ab} gives the result of the match between $a \in A$ and $b \in B$: it is 1 if a defeats b and 0 otherwise. Note that we do not allow for the possibility of draws, and every $a \in A$ faces every $b \in B$. When there is no ambiguity we denote a tournament simply by K , with the understanding that $A = [\text{rows}(K)]$ and $B = [\text{columns}(K)]$.

The *neighbourhood* of a player $a \in A$ in K is the set $K(a) = \{b \in B \mid K_{ab} = 1\} \subseteq B$, i.e. the set of players which a defeats. The




Figure 1: An example of a chain graph

neighbourhood of $b \in B$ is the set $K^{-1}(b) = \{a \in A \mid K_{ab} = 1\} \subseteq A$, i.e. the set of players defeating b .

Given a tournament K , our goal is to place a ranking on each of A and B . We define a ranking *operator* for this purpose.

Definition 2.2. An operator φ assigns each tournament K a pair $\varphi(K) = (\leq_K^\varphi, \sqsubseteq_K^\varphi)$ of total preorders on A and B respectively.²

For $a, a' \in A$, we interpret $a \leq_K^\varphi a'$ to mean that a' is ranked at least as strong as a in the tournament K , according to the operator φ (similarly, $b \sqsubseteq_K^\varphi b'$ means b' is ranked at least as strong as b). The strict and symmetric parts of \leq_K^φ are denoted by $<_K^\varphi$ and \approx_K^φ .

As a simple example, consider φ_{count} , where $a \leq_K^{\varphi_{\text{count}}} a'$ iff $|K(a)| \leq |K(a')|$ and $b \sqsubseteq_K^{\varphi_{\text{count}}} b'$ iff $|K^{-1}(b)| \geq |K^{-1}(b')|$. This operator simply ranks players by number of victories. It is a bipartite version of the *points system* introduced by Rubinstein [20], and generalises *Copeland’s rule* [5].

2.2 Chain Graphs

Each bipartite tournament K naturally corresponds to a bipartite graph G_K , with vertices $A \sqcup B$ and an edge between a and b whenever $K_{ab} = 1$.³ The task of ranking a tournament admits a particularly simple solution if this graph happens to be a *chain graph*.

Definition 2.3 ([26]). A bipartite graph $G = (U, V, E)$ is a *chain graph* if there is an ordering $U = \{u_1, \dots, u_k\}$ of U such that $N(u_1) \subseteq \dots \subseteq N(u_k)$, where $N(u_i) = \{v \in V \mid (u_i, v) \in E\}$ is the neighbourhood of u_i in G .

In other words, a chain graph is a bipartite graph where the neighbourhoods of the vertices on one side can be ordered so as to form a chain with respect to set inclusion. It is easily seen that this nesting property holds for U if and only if it holds for V . Figure 1 shows an example of a chain graph.

Now, as our terminology might suggest, the neighbourhood $K(a)$ of some player $a \in A$ in a tournament K coincides with the neighbourhood of the corresponding vertex in G_K . If G_K is a chain graph we can therefore enumerate A as $\{a_1, \dots, a_m\}$ such that $K(a_i) \subseteq K(a_{i+1})$ for each $1 \leq i < m$. This indicates that each a_{i+1} has performed *at least as well* as a_i in a strong sense: every opponent which a_i defeated was also defeated by a_{i+1} , and a_{i+1} may have additionally defeated opponents which a_i did not.⁴ It seems only

¹ Note that A and B are not disjoint as sets: 1 is always contained in both A and B , for instance. This poses no real problem, however, since we view the number 1 merely a *label* for a player. It will always be clear from context whether a given integer should be taken as a label for a player on the A side or the B side.

² A total preorder is a transitive and complete binary relation.

³ $A \sqcup B$ is the *disjoint union* of A and B , which we define as $\{(a, \mathcal{A}) \mid a \in A\} \cup \{(b, \mathcal{B}) \mid b \in B\}$, where \mathcal{A} and \mathcal{B} are constant symbols.

natural in this case that one should rank a_i (weakly) below a_{i+1} . Appealing to transitivity and the fact that each $a \in A$ appears as *some* a_i , we see that any tournament K where G_K is a chain graph comes pre-equipped with a natural total preorder on A , where a' ranks higher than a if and only if $K(a) \subseteq K(a')$. The duality of the neighbourhood-nesting property for chain graphs implies that B can also be totally preordered, with b' ranked higher than b if and only if $K^{-1}(b) \supseteq K^{-1}(b')$.⁵ Moreover, these total preorders relate to the tournament results in an important sense: if a defeats b and b' ranks worse than b , then a must defeat b' also. That is, the neighbourhood of each $a \in A$ is *downwards closed* w.r.t the ranking of B , and the neighbourhood of each $b \in B$ is *upwards closed* in A .

Tournaments corresponding to chain graphs will be said to satisfy the *chain property*, and will accordingly be called *chain tournaments*. We give a simpler (but equivalent) definition which does not refer to the underlying graph G_K . First, define relations $\leq_K^{\mathcal{A}}, \leq_K^{\mathcal{B}}$ on A and B respectively by $a \leq_K^{\mathcal{A}} a'$ iff $K(a) \subseteq K(a')$ and $b \leq_K^{\mathcal{B}} b'$ iff $K^{-1}(b) \supseteq K^{-1}(b')$, for any tournament K .

Definition 2.4. A tournament K has the *chain property* if $\leq_K^{\mathcal{A}}$ is a total preorder.

According to the duality principle mentioned already, the chain property implies that $\leq_K^{\mathcal{B}}$ is also a total preorder. Note that the relations $\leq_K^{\mathcal{A}}$ and $\leq_K^{\mathcal{B}}$ are analogues of the *covering relation* for non-bipartite tournaments [5].

Example 2.5. Consider $K = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix}$. Then $K(1) \subset K(2) \subset K(3)$, so K has the chain property. In fact, K is the tournament corresponding to the chain graph G from Figure 1.

3 RANKING VIA CHAIN EDITING

We have seen that chain tournaments come equipped with natural rankings of A and B . Such tournaments represent an ‘ideal’ situation, wherein the abilities of the players on both sides of the tournament are perfectly nested. Of course this may not be so in reality: the nesting may be broken by some $a \in A$ winning a match it ought not to by chance, or by losing a match by accident.

One idea for recovering a ranking in this case, originally suggested by Jiao et al. [15], is to apply *chain editing*: find the minimum number of edge changes required to convert the graph G_K into a chain graph. This process can be seen as correcting the ‘noise’ in an observed tournament K to obtain an ideal ranking. In this section we introduce the class of operators producing rankings in this way.

3.1 Chain-minimal Operators

To define chain-editing in our framework we once again present an equivalent definition which does not refer to the underlying graph G_K : the number of edge changes between graphs can be replaced by the *Hamming distance* between tournament matrices.

Definition 3.1. For $m, n \in \mathbb{N}$, let $C_{m,n}$ denote the set of all $m \times n$ chain tournaments. For an $m \times n$ tournament K , write $\mathcal{M}(K) =$

⁴ Note that this is a more robust notion of performance than comparing the neighbourhoods of a_i and a_{i+1} by *cardinality*, which may fail to account for differences in the strength of opponents when counting wins and losses.

⁵ Note that the ordering of the B s is reversed compared to the A s, since the larger $K^{-1}(b)$ the worse b has performed.

$\arg \min_{K' \in C_{m,n}} d(K, K') \subseteq \mathcal{K}$ for the set of chain tournaments closest to K w.r.t the Hamming distance $d(K, K') = |\{(a, b) \in A \times B \mid K_{ab} \neq K'_{ab}\}|$. Let $m(K)$ denote this minimum distance.

Note that chain editing, which is NP-hard in general [15], amounts to finding a single element of $\mathcal{M}(K)$.⁶ We comment further on the computational complexity of chain editing in Section 7. The following property characterises chain editing-based operators φ .

(chain-min) For every tournament K there is $K' \in \mathcal{M}(K)$ such that $\varphi(K) = (\leq_{K'}^{\mathcal{A}}, \leq_{K'}^{\mathcal{B}})$.

That is, the ranking of K is obtained by choosing the neighbourhood-subset rankings for some closest chain tournament K' . Operators satisfying **chain-min** will be called *chain-minimal*.

Example 3.2. Consider $K = \begin{bmatrix} 1 & 0 & 1 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 1 \end{bmatrix}$. K does not have the chain property, since neither $K(1) \subseteq K(2)$ nor $K(2) \subseteq K(1)$. The set $\mathcal{M}(K)$ consists of four tournaments a distance of 2 from K :

$$\mathcal{M}(K) = \left\{ \begin{bmatrix} 1 & 1 & 1 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 1 \end{bmatrix}, \begin{bmatrix} 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 1 \end{bmatrix}, \begin{bmatrix} 1 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 \end{bmatrix}, \begin{bmatrix} 1 & 0 & 1 & 0 \\ 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 \end{bmatrix} \right\}$$

The corresponding rankings are $(213, \{12\}34)$, $(123, 12\{34\})$, $(213, 13\{24\})$ and $(123, \{13\}24)$.⁷

Example 3.2 shows that there is no unique chain-minimal operator, since for a given tournament K there may be several closest chain tournaments to choose from. In Section 5 we introduce a principled way to single out a *unique* chain tournament and thereby construct a well-defined chain-minimal operator.

3.2 A Maximum Likelihood Interpretation

So far we have motivated **chain-min** as a way to fix errors in a tournament and recover the ideal or *true* ranking. In this section we make this notion precise by defining a probabilistic model in which chain-minimal rankings arise as maximum likelihood estimates. The maximum likelihood approach has been applied for (non-bipartite) tournaments (e.g. the Bradley-Terry model [4, 14]), voting in social choice theory [10], truth discovery [25], belief merging [11] and other related problems.

In this approach we take an epistemic view of tournament ranking: it is assumed there exists a true ‘state of the world’ which determines the tournament results along with objective rankings of A and B . A given tournament K is then seen as a *noisy observation* derived from the true state, and a *maximum likelihood estimate* is a state for which the probability of observing K is maximal.

More specifically, a state of the world is represented as a vector of *skill levels* for the players in A and B .⁸

Definition 3.3. For a fixed size $m \times n$, a *state of the world* is a tuple $\theta = \langle \mathbf{x}, \mathbf{y} \rangle$, where $\mathbf{x} \in \mathbb{R}^m$ and $\mathbf{y} \in \mathbb{R}^n$ satisfies the following properties:

$$\forall a, a' \in A \quad (x_a < x_{a'} \implies \exists b \in B : x_a < y_b \leq x_{a'}) \quad (1)$$

$$\forall b, b' \in B \quad (y_b < y_{b'} \implies \exists a \in A : y_b \leq x_a < y_{b'}) \quad (2)$$

where $A = [m]$, $B = [n]$. Write $\Theta_{m,n}$ for the set of all $m \times n$ states.

⁶ The decision problem associated with chain editing – which in tournament terms is the question of whether $m(K) \leq k$ for a given integer k – is NP-complete [9].

⁷ Here $a_1 a_2 a_3$ is shorthand for the ranking $a_1 < a_2 < a_3$ of A , and similar for B . Elements in brackets are ranked equally.

⁸ For simplicity we use numerical skill levels here, although it would suffice to have a partial preorder on $A \sqcup B$ such that each $a \in A$ is comparable with every $b \in B$.

For $a \in A$, x_a is the *skill level* of a in state θ (and similarly for y_b). These skill levels represent the true capabilities of the players in A and B in state θ : a is capable of defeating b if and only if $x_a \geq y_b$. Note that (1) suggests a simple form of *explainability*: a' can only be strictly more skilful than a if there is some $b \in B$ which *explains* this fact, i.e. some b which a' can defeat but a cannot ((2) is analogous for the B s). These conditions are intuitive if we assume that skill levels are relative to the sets A and B currently under consideration (i.e. they do not reflect the abilities of players in future matches against new contenders outside of A or B). Finally note that our states of the world are *richer* than the output of an operator, in contrast to other work in the literature [4, 10, 14]. Specifically, a state θ contains extra information in the form of comparisons between A and B .

Noise is introduced in the observed tournament K via *false positives* (where $a \in A$ defeats a more skilled $b \in B$ by accident) and *false negatives* (where $a \in A$ is defeated by an inferior $b \in B$ by mistake).⁹ The noise model is therefore parametrised by the false positive and false negative rates $\alpha = \langle \alpha_+, \alpha_- \rangle \in [0, 1]^2$, which we assume are the same for all $a \in A$.¹⁰ We also assume that noise occurs independently across all matches.

Definition 3.4. Let $\alpha = \langle \alpha_+, \alpha_- \rangle \in [0, 1]^2$. For each $m, n \in \mathbb{N}$ and $\theta = \langle \mathbf{x}, \mathbf{y} \rangle \in \Theta_{m,n}$, consider independent binary random variables X_{ab} representing the outcome of a match between $a \in [m]$ and $b \in [n]$, where

$$P_{\alpha}(X_{ab} = 1 \mid \theta) = \begin{cases} \alpha_+, & x_a < y_b \\ 1 - \alpha_-, & x_a \geq y_b \end{cases} \quad (3)$$

$$P_{\alpha}(X_{ab} = 0 \mid \theta) = \begin{cases} 1 - \alpha_+, & x_a < y_b \\ \alpha_-, & x_a \geq y_b \end{cases} \quad (4)$$

This defines a probability distribution $P_{\alpha}(\cdot \mid \theta)$ over $m \times n$ tournaments by

$$P_{\alpha}(K \mid \theta) = \prod_{(a,b) \in [m] \times [n]} P_{\alpha}(X_{ab} = K_{ab} \mid \theta)$$

Here $P_{\alpha}(K \mid \theta)$ is the probability of observing the tournament results K when the false positive and negative rates are given by α and the true state of the world is θ . Note that the four cases in (3) and (4) correspond to a false positive, true positive, true negative and false negative respectively. We can now define a maximum likelihood operator.

Definition 3.5. Let $\alpha \in [0, 1]^2$ and $m, n \in \mathbb{N}$. Then $\theta \in \Theta_{m,n}$ is a *maximum likelihood estimate* (MLE) for an $m \times n$ tournament K w.r.t α if $\theta \in \arg \max_{\theta' \in \Theta_{m,n}} P_{\alpha}(K \mid \theta')$. An operator φ is a *maximum likelihood operator* w.r.t α if for any $m, n \in \mathbb{N}$ and any $m \times n$ tournament K there is an MLE $\theta = \langle \mathbf{x}, \mathbf{y} \rangle \in \Theta_{m,n}$ for K such that $a \leq_K^{\varphi} a'$ iff $x_a \leq x_{a'}$ and $b \sqsubseteq_K^{\varphi} b'$ iff $y_b \leq y_{b'}$.

Now, consider the tournament K_{θ} associated with each state $\theta = \langle \mathbf{x}, \mathbf{y} \rangle$, given by $[K_{\theta}]_{ab} = 1$ if $x_a \geq y_b$ and $[K_{\theta}]_{ab} = 0$ otherwise. Note that K_{θ} is the unique tournament with non-zero probability when there are no false positive or false negatives. Expressed in terms of K_{θ} , the MLEs take a particularly simple form if $\alpha_+ = \alpha_-$, i.e. if false positives and false negatives occur at the same rate.

⁹Note that a false positive for a is a false negative for b and vice versa.

¹⁰This is a strong assumption, and it may be more realistic to model the false positive/negative rates as a function of x_a . We leave this to future work.

LEMMA 3.6. Let $\alpha = \langle \beta, \beta \rangle$ for some $\beta < \frac{1}{2}$. Then θ is an MLE for K if and only if $\theta \in \arg \min_{\theta' \in \Theta_{m,n}} d(K, K_{\theta'})$.

PROOF (SKETCH). Let K be an $m \times n$ tournament. It can be shown (and we do so in the appendix) that for any $\theta \in \Theta_{m,n}$

$$P_{\alpha}(K \mid \theta) = \left(\prod_{a \in A} \alpha_+^{|K(a) \setminus K_{\theta}(a)|} (1 - \alpha_-)^{|K(a) \cap K_{\theta}(a)|} \right) (1 - \alpha_+)^{|B \setminus (K(a) \cup K_{\theta}(a))|} \alpha_-^{|K_{\theta}(a) \setminus K(a)|}$$

Plugging in $\alpha_+ = \alpha_- = \beta$ and simplifying, one can obtain

$$P_{\alpha}(K \mid \theta) = c \prod_{a \in A} \left(\frac{\beta}{1 - \beta} \right)^{|K(a) \Delta K_{\theta}(a)|}$$

where $X \Delta Y = (X \setminus Y) \cup (Y \setminus X)$ is the symmetric difference of two sets X and Y , and $c = (1 - \beta)^{|A| \cdot |B|}$ is a positive constant that does not depend on θ . Now, $P_{\alpha}(K \mid \theta)$ is positive, and is maximal when its logarithm is. We have

$$\begin{aligned} \log P_{\alpha}(K \mid \theta) &= \log c + \log \left(\frac{\beta}{1 - \beta} \right) \sum_{a \in A} |K(a) \Delta K_{\theta}(a)| \\ &= \log c + \log \left(\frac{\beta}{1 - \beta} \right) d(K, K_{\theta}) \end{aligned}$$

Since $\log c$ is constant and $\beta < 1/2$ implies $\log \left(\frac{\beta}{1 - \beta} \right) < 0$, it follows that $\log P_{\alpha}(K \mid \theta)$ is maximised exactly when $d(K, K_{\theta})$ is minimised, which proves the result. \square

This result characterises the MLE states for K as those for which K_{θ} is the closest to K . As it turns out, the tournaments K_{θ} that arise in this way are exactly those with the chain property.

LEMMA 3.7. An $m \times n$ tournament K has the chain property if and only if $K = K_{\theta}$ for some $\theta \in \Theta_{m,n}$.

The proof of Lemma 3.7 relies crucially on (1) and (2) in the definition of a state. Combining all the results so far we obtain our first main result: the maximum likelihood operators for $\alpha = \langle \beta, \beta \rangle$ are exactly the chain-minimal operators.

THEOREM 3.8. Let $\alpha = \langle \beta, \beta \rangle$ for some $\beta < \frac{1}{2}$. Then φ is a maximum likelihood operator w.r.t α if and only if φ satisfies **chain-min**.

PROOF (SKETCH). First note that by Lemma 3.6, a state θ is an MLE for an $m \times n$ tournament K iff K_{θ} is closest to K amongst all other tournaments $\{K_{\theta'} \mid \theta' \in \Theta_{m,n}\}$. But by Lemma 3.7, this set is exactly the $m \times n$ tournaments with the chain property. It follows from the definition of $\mathcal{M}(K)$ that θ is an MLE if and only if $K_{\theta} \in \mathcal{M}(K)$. Consequently, $K' \in \mathcal{M}(K)$ if and only if $K' = K_{\theta}$ for some MLE θ for K . We see that **chain-min** can be equivalently stated as follows: for all K there exists an MLE θ such that $\varphi(K) = (\leq_{K_{\theta}}^{\mathcal{A}}, \leq_{K_{\theta}}^{\mathcal{B}})$. Using properties (1) and (2) in Definition 3.3 for θ it is straightforward to show that $a \leq_{K_{\theta}}^{\mathcal{A}} a'$ iff $x_a \leq x_{a'}$ and $b \leq_{K_{\theta}}^{\mathcal{B}} b'$ iff $y_b \leq y_{b'}$ for all $a, a' \in A, b, b' \in B$ (where $\theta = \langle \mathbf{x}, \mathbf{y} \rangle$). This means that the above reformulation of **chain-min** coincides with the definition of a maximum likelihood operator, and we are done. \square

Similar results can be obtained for other limiting values of α . If $\alpha_+ = 0$ and $\alpha_- \in (0, 1)$ then the MLE operators correspond to *chain completion*: finding the minimum number of edge additions required

to make G_K a chain graph. This models situations where false positives never occur, although false negatives may (e.g. numerical entry questions in the case where A represents students and B exam questions [15]). Similarly, the case $\alpha_- = 0$ and $\alpha_+ \in (0, 1)$ corresponds to *chain deletion*, where edge additions are not allowed.

4 AXIOMATIC ANALYSIS

Chain-minimal operators have theoretical backing in a probabilistic sense due to the results of Section 3.2, but are they appropriate ranking methods in practise? To address this question we consider the *normative* properties of chain-minimal operators via the axiomatic method of social choice theory. We formulate several axioms for bipartite tournament ranking and assess whether they are compatible with **chain-min**. It will be seen that an important *anonymity* axiom fails for all chain-minimal operators; later in Section 5 we describe a scenario in which this is acceptable and define a class of concrete operators for this case, and in Section 6 we relax the **chain-min** requirement in order to gain anonymity.

4.1 The Axioms

We will consider five axioms – mainly adaptations of standard social choice properties to the bipartite tournament setting.

Symmetry Properties. We consider two symmetry properties. The first is a classic *anonymity* axiom, which says that an operator φ should not be sensitive to the ‘labels’ used to identify participants in a tournament. Axioms of this form are standard in social choice theory; a tournament version goes at least as far back as [20].

We need some notation: for a tournament K and permutations $\sigma : A \rightarrow A$, $\pi : B \rightarrow B$, let $\sigma(K)$ and $\pi(K)$ denote the tournament obtained by permuting the rows and columns of K by σ and π respectively, i.e. $[\sigma(K)]_{ab} = K_{\sigma^{-1}(a), b}$ and $[\pi(K)]_{ab} = K_{a, \pi^{-1}(b)}$. Note that in the statement of the axioms we omit universal quantification over K , $a, a' \in A$ and $b, b' \in B$ for brevity.

(anon) Let $\sigma : A \rightarrow A$ and $\pi : B \rightarrow B$ be permutations. Then $a \leq_K^\varphi a'$ iff $\sigma(a) \leq_{\pi(\sigma(K))}^\varphi \sigma(a')$.

Our second axiom is specific to bipartite tournaments, and expresses a *duality* between the two sides A and B : given the two sets of conceptually disjoint entities participating in a bipartite tournament, it should not matter which one we label A and which one we label B . We need the notion of a *dual tournament*.

Definition 4.1. The *dual tournament* of K is $\bar{K} = \mathbf{1} - K^\top$, where $\mathbf{1}$ denotes the matrix consisting entirely of 1s.

\bar{K} is essentially the same tournament as K , but with the roles of A and B swapped. In particular, $A_{\bar{K}} = B_K$, $B_{\bar{K}} = A_K$ and $K_{ab} = 1$ iff $\bar{K}_{ba} = 0$. Also note that $\bar{\bar{K}} = K$. The duality axiom states that the ranking of the B s in K is the same as the A s in \bar{K} .

(dual) $b \sqsubseteq_K^\varphi b'$ iff $b \leq_{\bar{K}}^\varphi b'$.

Whilst **dual** is not necessarily a universally desirable property – one can imagine situations where A and B are not fully abstract and should not be treated symmetrically – it is important to consider in any study of bipartite tournaments. Note that **dual** implies $a \leq_K^\varphi a'$ iff $a \sqsubseteq_{\bar{K}}^\varphi a'$, so that a **dual**-operator can be defined by giving the ranking for one of A or B only, and defining the other by

duality. This explains our choice to define **anon** (and subsequent axioms) solely in terms of the A ranking: the analogous anonymity constraint for the B ranking follows from **anon** together with **dual**.

An Independence Property. *Independence axioms* play a crucial role in social choice. We present a bipartite adaptation of a classic axiom introduced in [20], which has subsequently been called *Independence of Irrelevant Matches* [14].

(IIM) If K_1, K_2 are tournaments of the same size with identical a -th and a' -th rows, then $a \leq_{K_1}^\varphi a'$ iff $a \leq_{K_2}^\varphi a'$.

IIM is a strong property, which says the relative ranking of a and a' does not depend on the results of any match not involving a or a' . This axiom has been questioned for generalised tournaments [14], and a similar argument can be made against it here: although each player in A faces the same opponents, we may wish to take the *strength* of opponents into account, e.g. by rewarding victories against highly-ranked players in B . Consequently we do not view **IIM** as an essential requirement, but rather introduce it to facilitate comparison with our work and the existing tournament literature.

Monotonicity Properties. Our final axioms are monotonicity properties, which express the idea that *more victories are better*. The first axiom follows our original intuition for constructing the natural ranking associated with a chain graph; namely that $K(a) \subseteq K(a')$ indicates a' has performed at least as well as a .

(mon) If $K(a) \subseteq K(a')$ then $a \leq_K^\varphi a'$.

Note that **mon** simply says \leq_K^φ extends the (in general, partial) preorder $\leq_K^{\mathcal{A}}$. Yet another standard axiom is *positive responsiveness*.

(pos-resp) If $a \leq_K^\varphi a'$ and $K_{a',b} = 0$ for some $b \in B$, then $a <_{K+1_{a',b}}^\varphi a'$, where $1_{a',b}$ is the matrix with 1 in position (a', b) and zeros elsewhere.

That is, adding an extra victory for a should only improve its ranking, with ties now broken in its favour. This version of positive responsiveness was again introduced in [20], where together with **anon** and **IIM** it characterises the *points system* ranking method for round-robin tournaments, which simply ranks players according to the number of victories. The analogous operator in our framework is φ_{count} , and it can be shown that φ_{count} is uniquely characterised by **anon**, **IIM**, **pos-resp** and **dual**. Finally, note that **pos-resp** also acts as a kind of *strategyproofness*: a cannot improve its ranking by deliberately losing a match. Specifically, if $K_{ab} = 1$ and $a \leq_K^\varphi a'$, then **pos-resp** implies $a <_{K-1_{ab}}^\varphi a'$.

4.2 Axiom Compatibility with chain-min

We come to analysing the compatibility of **chain-min** with the axioms. First, the negative results.

THEOREM 4.2. *There is no operator satisfying **chain-min** and any of **anon**, **IIM** or **pos-resp**.*

The counterexample for **anon** is particularly simple: take $K = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$. Swapping the rows and columns brings us back to K , so **anon** implies 1, 2 $\in A$ rank equally. However, it is easily seen for every $K' \in \mathcal{M}(K)$, either $K(1) \subset K(2)$ or $K(2) \subset K(1)$, i.e. no chain-minimal operator can rank 1 and 2 equally.

The MLE results of Section 3.2 provides informal explanation for this result. For K above to arise in the noise model of Definition 3.4 there must have been two ‘mistakes’ (false positives or false negatives). This is less likely than a single mistake from just one of $1, 2 \in A$, but the likelihood maximisation forces us to choose one or the other. A similar argument explains the **pos-resp** failure.

It is also worth noting that **anon** only fails at the last step of chain editing, where a single element of $\mathcal{M}(K)$ is chosen. Indeed, the set $\mathcal{M}(K)$ itself *does* exhibit the kind of symmetry one might expect: we have $\mathcal{M}(\pi(\sigma(K))) = \{\pi(\sigma(K')) \mid K' \in \mathcal{M}(K)\}$. This means that an operator which aggregates the rankings from *all* $K' \in \mathcal{M}(K)$ – e.g. any anonymous social welfare function – would satisfy **anon**. The other axioms are compatible with **chain-min**.

THEOREM 4.3. *For each of **dual** and **mon**, there exists an operator satisfying **chain-min** and the stated property.*

Despite the simplicity of **mon**, Theorem 4.3 is deceptively difficult to prove. We describe operators satisfying **chain-def** and **dual** or **mon** non-constructively by first taking an *arbitrary* chain-minimal operator φ , and using properties of the set $\mathcal{M}(K)$ to produce φ' satisfying **dual** or **mon**. Note also that we have not yet constructed an operator satisfying **dual**, **mon** and **chain-min** simultaneously, although we conjecture that such operators do exist.

5 MATCH-PREFERENCE OPERATORS

The counterexample for **chain-min** and **anon** suggests that chain-minimal operators require some form of tie-breaking mechanism when the tournaments in $\mathcal{M}(K)$ cannot be distinguished while respecting anonymity. While this limits the use of chain-minimal operators as general purpose ranking methods, it is not such a problem if additional information is available to guide the tie-breaking. In this section we introduce a new class of operators for this case.

The core idea is to single out a unique chain tournament close to K by paying attention to not only the *number* of entries in K that need to be changed to produce a chain tournament, *which* entries. Specifically, we assume the availability of a total order on the set of matrix indices $\mathbb{N} \times \mathbb{N}$ (the *matches*) which indicates our willingness to change an entry in K : the higher up (a, b) is in the ranking, the more acceptable it is to change K_{ab} during chain editing.

This total order – called the *match-preference relation* – is fixed for all tournaments K ; this means we are dealing with extra information about how tournaments are *constructed in matrix form*, not extra information about any specific tournament K .

One possible motivation for such a ranking comes from cases where matches occur at distinct points in time. In this case the matches occurring more recently are (presumably) more representative of the players’ *current* abilities, and we should therefore prefer to modify the outcome of old matches where possible.

For the formal definition we need notation for the *vectorisation* of a tournament K : for a total order \preceq on $\mathbb{N} \times \mathbb{N}$ and an $m \times n$ tournament K , we write $\text{vec}_{\preceq}(K)$ for the vector in $\{0, 1\}^{mn}$ obtained by collecting the entries of K in the order given by $\preceq \upharpoonright (A \times B)$,¹¹ starting with the minimal entry. That is, $\text{vec}_{\preceq}(K) = (K_{a_1, b_1}, \dots, K_{a_{mn}, b_{mn}})$, where $(a_1, b_1), \dots, (a_{mn}, b_{mn})$ is the unique enumeration of $A \times B$ such that $(a_i, b_i) \preceq (a_{i+1}, b_{i+1})$ for each i .

¹¹ This denotes the restriction of \preceq to $A \times B$, i.e. $\preceq \cap ((A \times B) \times (A \times B))$.

The operator corresponding to \preceq is defined using the notion of a *choice function*: a function α which maps any tournament K to an element of $\mathcal{M}(K)$. Any such function defines a chain-minimal operator φ by setting $\varphi(K) = (\preceq_{\alpha(K)}^{\mathcal{A}}, \preceq_{\alpha(K)}^{\mathcal{B}})$.

Definition 5.1. Let \preceq be a total order on $\mathbb{N} \times \mathbb{N}$. Define an operator φ_{\preceq} according to the choice function

$$\alpha_{\preceq}(K) = \arg \min_{K' \in \mathcal{M}(K)} \text{vec}_{\preceq}(K \oplus K') \quad (5)$$

where $[K \oplus K']_{ab} = |K_{ab} - K'_{ab}|$, and the minimum is taken w.r.t the lexicographic ordering on $\{0, 1\}^{|A| \cdot |B|}$.¹² Operators generated in this way will be called *match-preference operators*.

Example 5.2. Let \preceq be the lexicographic order¹³ on $\mathbb{N} \times \mathbb{N}$ so that $\text{vec}_{\preceq}(K \oplus K')$ is obtained by collecting the entries of $K \oplus K'$ row-by-row, from top to bottom and left to right. Take K from Example 3.2. Writing K_1, \dots, K_4 for the elements of $\mathcal{M}(K)$ in the order that they appear in Example 3.2 and setting $v_i = \text{vec}_{\preceq}(K \oplus K_i)$, we have

$$\begin{aligned} v_1 &= (0100 \ 0000 \ \mathbf{10000}); & v_2 &= (0010 \ 0000 \ \mathbf{10000}) \\ v_3 &= (0000 \ \mathbf{0100} \ \mathbf{10000}); & v_4 &= (0000 \ 0010 \ \mathbf{10000}) \end{aligned}$$

The lexicographic minimum is the one with the 1 entries as far right as possible, which in this case is v_4 . Consequently φ_{\preceq} ranks K according to K_4 , i.e. $1 \prec_K^{\varphi_{\preceq}} 2 \prec_K^{\varphi_{\preceq}} 3$ and $1 \approx_K^{\varphi_{\preceq}} 3 \sqsubset_K^{\varphi_{\preceq}} 2 \sqsubset_K^{\varphi_{\preceq}} 4$.

To conclude the discussion of match-preference operators, we note that one can compute $\alpha_{\preceq}(K)$ as the unique closest chain tournament to K w.r.t a *weighted* Hamming distance, and thereby avoid the need to enumerate $\mathcal{M}(K)$ in full as per eq. (5).

THEOREM 5.3. *Let \preceq be a total order on $\mathbb{N} \times \mathbb{N}$. Then for any $m, n \in \mathbb{N}$ there exists a function $w : [m] \times [n] \rightarrow \mathbb{R}_{\geq 0}$ such that for all $m \times n$ tournaments K :*

$$\arg \min_{K' \in \mathcal{C}_{m,n}} d_w(K, K') = \{\alpha_{\preceq}(K)\} \quad (6)$$

where $d_w(K, K') = \sum_{(a,b) \in [m] \times [n]} w(a,b) \cdot |K_{ab} - K'_{ab}|$.

For example, the weights corresponding to \preceq from Example 5.2 and $m = 2, n = 3$ are $w = \begin{bmatrix} 1.5 & 1.25 & 1.125 \\ 1.0625 & 1.03125 & 1.015625 \end{bmatrix}$.

6 RELAXING CHAIN-MIN

Having studied chain-minimal operators in some detail, we turn to two remaining problems: **chain-min** is incompatible with **anon**, and computing a chain-minimal operator is NP-hard. In this section we obtain both anonymity and tractability by relaxing the **chain-min** requirement to a property we call *chain-definability*. We go on to characterise the class of operators with this weaker property via a greedy approximation algorithm, single out a particularly intuitive instance, and revisit the axioms of Section 4.

6.1 Chain-definability

The source of the difficulties with **chain-min** lies in the minimisation aspect of chain editing. A natural way to retain the spirit of **chain-min** without the complications is to require that $\varphi(K)$ corresponds to *some* chain tournament, not necessarily one closest to K . We call this property *chain-definability*.

¹² Note that $K \oplus K'$ is 1 in exactly the entries where K and K' differ.

¹³ That is, $(a, b) \preceq (a', b')$ iff $a < a'$ or $(a = a'$ and $b \leq b')$.

(chain-def) For every $m \times n$ tournament K there is $K' \in C_{m,n}$ such that $\varphi(K) = (\leq_{K'}^{\mathcal{A}}, \leq_{K'}^{\mathcal{B}})$.

Clearly **chain-min** implies **chain-def**. ‘Chain-definable’ operators can also be cast in the MLE framework of Section 3.2 as those whose rankings correspond to *some* (not necessarily MLE) state θ .

At first glance it may seem difficult to determine whether a given pair of rankings correspond to a chain tournament, since the number of such tournaments grows rapidly with m and n . Fortunately, **chain-def** can be characterised without reference to chain tournaments by considering the number of *ranks* of \leq_K^φ and \sqsubseteq_K^φ . In what follows $\text{ranks}(\leq)$ denotes the number of ranks of a total preorder \leq , i.e. the number of equivalence classes of its symmetric part.

THEOREM 6.1. φ satisfies **chain-def** if and only if $|\text{ranks}(\leq_K^\varphi) - \text{ranks}(\sqsubseteq_K^\varphi)| \leq 1$ for every tournament K .

6.2 Interleaving Operators

According to Theorem 6.1, to construct a chain-definable operator it is enough to ensure that the number of ranks of \leq_K^φ and \sqsubseteq_K^φ differ by at most one. A simple way to achieve this is to iteratively select and remove the top-ranked players of A and B simultaneously, until one of A or B is exhausted. We call such operators *interleaving operators*. Closely related ranking methods have been previously introduced for non-bipartite tournaments by Bouyssou [2].

Formally, our procedure is defined by two functions f and g which select the next top ranks given a tournament K and subsets $A' \subseteq A$, $B' \subseteq B$ of the remaining players.

Definition 6.2. An \mathcal{A} -selection function is a mapping $f : \mathcal{K} \times 2^{\mathcal{N}} \times 2^{\mathcal{N}} \rightarrow 2^{\mathcal{N}}$ such that for any tournament K , $A' \subseteq A$ and $B' \subseteq B$: (i) $f(K, A', B') \subseteq A'$; (ii) If $A' \neq \emptyset$ then $f(K, A', B') \neq \emptyset$; (iii) $f(K, A', \emptyset) = A'$.

Similarly, a \mathcal{B} -selection function is a mapping $g : \mathcal{K} \times 2^{\mathcal{N}} \times 2^{\mathcal{N}} \rightarrow 2^{\mathcal{N}}$ such that (i) $g(K, A', B') \subseteq B'$; (ii) If $B' \neq \emptyset$ then $g(K, A', B') \neq \emptyset$; (iii) $g(K, \emptyset, B') = B'$.

The corresponding interleaving operator ranks players according to how soon they are selected in this way; the earlier the better.

Definition 6.3. Let f and g be selection functions and K a tournament. Write $A_0 = A$, $B_0 = B$, and for $i \geq 0$:

$$A_{i+1} = A_i \setminus f(K, A_i, B_i); \quad B_{i+1} = B_i \setminus g(K, A_i, B_i)$$

For $a \in A$ and $b \in B$, write $r(a) = \max\{i \mid a \in A_i\}$ and $s(b) = \max\{i \mid b \in B_i\}$.¹⁴ We define the corresponding *interleaving operator* $\varphi = \varphi_{f,g}^{\text{int}}$ by $a \leq_K^\varphi a'$ iff $r(a) \geq r(a')$ and $b \sqsubseteq_K^\varphi b'$ iff $s(b) \geq s(b')$.

Note that A_i and B_i are the players left remaining after i applications of f and g , i.e. after removing the top i ranks from both sides. Before giving a concrete example, we note that interleaving is not just *one* way to satisfying **chain-def**, it is the *only* way.

THEOREM 6.4. An operator φ satisfies **chain-def** if and only if $\varphi = \varphi_{f,g}^{\text{int}}$ for some selection functions (f, g) .

¹⁴ We show in the appendix that the recursive procedure eventually terminates with A_i and B_i becoming empty (and remaining so) after finitely many iterations, so r and s are well-defined.

Table 1: Iteration of the interleaving algorithm for φ_{CI}

i	K	A_i	B_i	f	g	K'_i
0	$\begin{bmatrix} 1 & 1 & 1 & 1 & 0 \\ 0 & 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 1 & 1 \\ 0 & 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 1 & 0 \end{bmatrix}$	{1, 2, 3, 4}	{1, 2, 3, 4, 5}	{1}	{1}	$\begin{bmatrix} 1 & 1 & 1 & 1 & 1 \\ 0 & 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 1 & 1 \\ 0 & 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 \end{bmatrix}$
1	$\begin{bmatrix} 0 & 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 1 & 1 \\ 0 & 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 1 & 0 \end{bmatrix}$	{2, 3, 4}	{2, 3, 4, 5}	{3}	{3, 4}	$\begin{bmatrix} 0 & 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 1 & 1 \\ 0 & 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 \end{bmatrix}$
2	$\begin{bmatrix} 0 & 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 1 & 1 \\ 0 & 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 1 & 0 \end{bmatrix}$	{2, 4}	{2, 5}	{2}	{5}	-
3	$\begin{bmatrix} 0 & 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 1 & 1 \\ 0 & 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 1 & 0 \end{bmatrix}$	{4}	{2}	{4}	{2}	-
4	-	\emptyset	\emptyset	\emptyset	\emptyset	-

Theorem 6.4 justifies our study of interleaving operators, and provides a different perspective on chain-definability via the selection functions f and g . We come to an important example.

Example 6.5. Define the *cardinality-based interleaving operator* $\varphi_{\text{CI}} = \varphi_{f,g}^{\text{int}}$ where $f(K, A', B') = \arg \max_{a \in A'} |K(a) \cap B'|$ and $g(K, A', B') = \arg \min_{b \in B'} |K^{-1}(b) \cap A'|$, so that the ‘winners’ at each iteration are the A s with the most wins, and the B s with the least losses, when restricting to A' and B' only. We take the $\arg \min/\arg \max$ to be the emptyset whenever A' or B' is empty.

Table 1 shows the iteration of the algorithm for a 4×5 tournament K . In each row i we show K with the rows and columns of $A \setminus A_i$ and $B \setminus B_i$ greyed out, so as to make it more clear how the f and g values are calculated.¹⁵ For brevity we also write f and g in place of $f(K, A_i, B_i)$ and $g(K, A_i, B_i)$ respectively.

The r and s values can be read off as 0, 2, 1, 3 for A and 0, 3, 1, 1, 2 for B , giving the ranking on A as $4 < 2 < 3 < 1$, and the ranking on B as $2 \sqsubset 5 \sqsubset 3 \approx 4 \sqsubset 1$. Note also that each $f(K, A_i, B_i)$ is a rank of \leq_K^φ (and similar for $g(K, A_i, B_i)$), so the rankings can in fact be read off by looking at the f and g columns of Table 1.

The interleaving algorithm can also be seen as a greedy algorithm for converting K into a chain graph directly. Indeed, by setting the neighbourhood of each $a \in f(K, A_i, B_i)$ to B_i , and removing each $b \in g(K, A_i, B_i)$ from the neighbourhoods of all $a \in A_{i+1}$, we eventually obtain a chain graph. We show this process in the K'_i column of Table 1, where only three entries need to be changed.¹⁶ The selection functions f and g can therefore be seen as *heuristics* with the goal of finding a chain graph ‘close’ to K .

The operator φ_{CI} from Example 6.5 uses simple cardinality-based heuristics, and can be seen as a chain-definable version of φ_{count} (which is not chain-definable). It is also the bipartite counterpart to repeated applications of Copeland’s rule [2]. Note that $f(K, A_i, B_i)$ and $g(K, A_i, B_i)$ can be computed in $O(N^2)$ time at each iteration i , where $N = |A| + |B|$. Since there cannot be more than N iterations, it follows that the rankings of φ_{CI} can be computed in $O(N^3)$ time.

6.3 Axiom Compatibility

We now revisit the axioms of Section 4 in relation to chain-definable operators in general and φ_{CI} specifically. Firstly, the weakening of **chain-min** pays off: **chain-def** is compatible with all our axioms.

¹⁵ Note that while f and g for φ_{CI} are independent of the greyed out entries, we do not require this property for selection functions in general.

¹⁶ In this example $\mathcal{M}(K)$ contains a single tournament a distance of 2 from K , so φ_{CI} makes one more change than necessary.

THEOREM 6.6. *For each of **anon**, **dual**, **IIM**, **mon** and **pos-resp**, there exists an operator satisfying **chain-def** and the stated property.*

Unfortunately, these cannot all hold at the same time. Indeed, taking $K = \begin{bmatrix} 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 \end{bmatrix}^T$ and assuming **anon** and **pos-resp**, the ranking on A is fully determined as $1 < 2 \approx 3 < 4$, and $\text{ranks}(\leq_K^\varphi) = 3$. However, **anon** with **dual** implies the ranking of B is flat, i.e. $\text{ranks}(\leq_K^\varphi) = 1$. This contradicts **chain-def** by Theorem 6.1, yielding the following impossibility result.

THEOREM 6.7. *There is no operator satisfying **chain-def**, **anon**, **dual** and **pos-resp**.*

For the specific operator φ_{C1} we have the following.

THEOREM 6.8. *φ_{C1} satisfies **chain-def**, **anon**, **dual** and **mon**, and does not satisfy **IIM** or **pos-resp**.*

Note that **anon** is satisfied. This makes φ_{C1} an important example of a well-motivated, tractable, chain-definable and anonymous operator, meeting the criteria outlined at the start of this section.

7 RELATED WORK

On chain graphs. Chain graphs were originally introduced by Yannakakis [26], who proved that *chain completion* – finding the minimum number of edges that when added to a bipartite graph form a chain graph – is NP-complete. Hardness results have subsequently been obtained for chain *deletion* [18] (where only edge deletions are allowed) and chain *editing* [9] (where both additions and deletions are allowed). We refer the reader to the work of Jiao et al. [15] and Drange et al. [9] for a more detailed account of this literature. Outside of complexity theory, chain graphs have been studied for their spectral properties in [1, 13], and the more general notion of a *nested colouring* was introduced in [6].

On tournaments in social choice. Tournaments have important applications in the design of voting rules, where an alternative x beats y in a pairwise comparison if a majority of voters prefer x to y . Various *tournament solutions* have been proposed, which select a set of ‘winners’ from a given tournament.¹⁷ Of particular relevance to our work are the *Slater set* and *Kemeney’s rule* [5], which find minimal sets of edges to invert in the tournament graph such that the beating relation becomes a total order.¹⁸ These methods are intuitively similar to chain editing: both involve making minimal changes to the tournament until some property is satisfied. A rough analogue to the Slater set in our framework is the union of the top-ranked players from each $K' \in \mathcal{M}(K)$. Solutions based on the covering relation – such as the *uncovered* and *Banks set* [5] – also bear similarity to chain editing.

Finally, note that directed versions of chain graphs (obtained by orienting edges from A to B and adding missing edges from B to A) correspond to *acyclic tournaments*, and a topological sort of A becomes a linearisation of the chain ranking \leq_K^A . This suggests a connection between chain deletion and the standard *feedback arc set* problem for removing cycles and obtaining a ranking.

¹⁷Note that a ranking, such as we consider in this paper, induces a set of winners by taking the maximally ranked players.

¹⁸Note that like chain editing, Kemeny’s rule also admits a maximum likelihood characterisation [10].

On generalised tournaments. A *generalised tournament* [14] is a pair (X, T) , where $X = [t]$ for some $t \in \mathbb{N}$ and $T \in \mathbb{R}_{\geq 0}^{t \times t}$ is a non-negative $t \times t$ matrix with $T_{ii} = 0$ for all $i \in X$. In this formalism each encounter between a pair of players i and j is represented by two numbers: T_{ij} and T_{ji} . This allows one to model both intensities of victories and losses (including draws) via the difference $T_{ij} - T_{ji}$, and the case where a comparison is not available (where $T_{ij} = T_{ji} = 0$).

Any $m \times n$ bipartite tournament K has a natural generalised tournament representation via the $(m+n) \times (m+n)$ *anti-diagonal block matrix* $T = \begin{bmatrix} 0 & K \\ K & 0 \end{bmatrix}$, where the top-left and bottom-right blocks are the $m \times m$ and $n \times n$ zero matrices respectively. However, such anti-diagonal block matrices are often excluded in the generalised tournament literature due to an assumption of *irreducibility*, which requires that the directed graph corresponding to T is strongly connected. This is not the case in general for T constructed as above, which means not all existing tournament operators (and tournament axioms) are well-defined for bipartite inputs.¹⁹ Consequently, bipartite tournaments are a special case of generalised tournaments *in principle*, but not in practise.

8 CONCLUSION

Summary. In this paper we studied chain editing, an interesting problem from computational complexity theory, as a ranking mechanism for bipartite tournaments. We analysed such mechanisms from a probabilistic viewpoint via the MLE characterisation, and in axiomatic terms. To resolve both the failure of an important anonymity axiom and NP-hardness, we weakened the chain editing requirement to one of *chain definability*, and characterised the resulting class of operators by the intuitive interleaving algorithm.

Limitations and future work. The hardness of chain editing remains a limitation of our approach. A possible remedy is to look to one of the numerous variant problems that are polynomial-time solvable [15]; determining their applicability to ranking is an interesting topic for future work. One could develop approximation algorithms for chain editing, possibly based on existing approximations of chain completion [17]. The interleaving operators of Section 6.2 go in this direction, but we did not yet obtain any theoretical or experimental bounds on the approximation ratio.

A second limitation of our work lies in the assumptions of the probabilistic model; namely that the true state of the world can be reduced to vectors of numerical skill levels which totally describe the tournament participants. This assumption may be violated when the competitive element of a tournament is *multi-faceted*, since a single number cannot represent multiple orthogonal components of a player’s capabilities. Nevertheless, if skill levels are taken as *aggregations* of these components, chain editing may prove to be a useful, albeit simplified, model.

Finally, there is room for more detailed axiomatic investigation. In this paper we have stuck with fairly standard social choice axioms and performed preliminary analysis. However, the indirect nature of the comparisons in a bipartite tournament presents unique challenges; new axioms may need to be formulated to properly evaluate bipartite ranking methods in a normative sense.

¹⁹We note that Slutzki and Volij [23] side-step the reducibility issue by decomposing T into irreducible components and ranking each separately, although their methods may give only *partial orders*.

REFERENCES

- [1] Milica Andelić, Enide Andrade, Domingos M. Cardoso, Carlos M. da Fonseca, Slobodan K. Simić, and Dejan V. Tošić. 2015. Some new considerations about double nested graphs. *Linear Algebra Appl.* 483 (2015), 323 – 341.
- [2] Denis Bouyssou. 2004. Monotonicity of ‘ranking by choosing’: A progress report. *Social Choice and Welfare* 23, 2 (2004), 249–273.
- [3] Sándor Bozóki, László Csató, and József Temesi. 2016. An application of incomplete pairwise comparison matrices for ranking top tennis players. *European Journal of Operational Research* 248, 1 (2016), 211–218.
- [4] Ralph Allan Bradley and Milton E. Terry. 1952. Rank Analysis of Incomplete Block Designs: I. The Method of Paired Comparisons. *Biometrika* 39, 3/4 (1952), 324–345.
- [5] Felix Brandt, Markus Brill, Paul Harrenstein, and Hervé Moulin. 2016. Tournament Solutions. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press, 57–84.
- [6] David Cook II. 2015. Nested colorings of graphs. *Australasian Journal of Combinatorics* 62, 1 (2015), 100–127.
- [7] László Csató. 2019. An impossibility theorem for paired comparisons. *Central European Journal of Operations Research* 27, 2 (2019), 497–514.
- [8] Paul Denny, John Hamer, Andrew Luxton-Reilly, and Helen Purchase. 2008. PeerWise: Students Sharing Their Multiple Choice Questions. In *Proceedings of the Fourth International Workshop on Computing Education Research (ICER '08)*. ACM, New York, 51–58.
- [9] Pål Grønås Drange, Markus Sortland Dregi, Daniel Lokshtanov, and Blair D. Sullivan. 2015. On the threshold of intractability. In *Algorithms-ESA 2015*. Springer, 411–423.
- [10] Edith Elkind and Arkadii Slinko. 2016. Rationalizations of voting rules. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press, 169–196.
- [11] Patricia Everaere, Sebastien Konieczny, and Pierre Marquis. 2020. Belief Merging Operators as Maximum Likelihood Estimators. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, Christian Bessiere (Ed.). International Joint Conferences on Artificial Intelligence Organization, 1763–1769.
- [12] Alban Galland, Serge Abiteboul, Amélie Marian, and Pierre Senellart. 2010. Corroborating Information from Disagreeing Views. In *Proceedings of the Third ACM International Conference on Web Search and Data Mining (WSDM '10)*. ACM, New York, 131–140.
- [13] Ebrahim Ghorbani. 2017. Some spectral properties of chain graphs. (2017). arXiv:math.CO/1703.03581
- [14] Julio González-Díaz, Ruud Hendrickx, and Edwin Lohmann. 2014. Paired comparisons analysis: an axiomatic approach to ranking methods. *Social Choice and Welfare* 42, 1 (2014), 139–169.
- [15] Yang Jiao, R Ravi, and Wolfgang Gatterbauer. 2017. Algorithms for automatic ranking of participants and tasks in an anonymized contest. In *International Workshop on Algorithms and Computation*. Springer, 335–346.
- [16] Yaliang Li, Jing Gao, Chuishi Meng, Qi Li, Lu Su, Bo Zhao, Wei Fan, and Jiawei Han. 2016. A Survey on Truth Discovery. *SIGKDD Explor. Newsl.* 17, 2 (Feb. 2016), 1–16.
- [17] Assaf Natanzon, Ron Shamir, and Roded Sharan. 2000. A polynomial approximation algorithm for the minimum fill-in problem. *SIAM J. Comput.* 30, 4 (2000), 1067–1079.
- [18] Assaf Natanzon, Ron Shamir, and Roded Sharan. 2001. Complexity classification of some edge modification problems. *Discrete Applied Mathematics* 113, 1 (2001), 109–128.
- [19] Catherine Olsson, Surya Bhupatiraju, Tom Brown, Augustus Odena, and Ian Goodfellow. 2018. Skill Rating for Generative Models. (2018). arXiv:stat.ML/1808.04888
- [20] Ariel Rubinstein. 1980. Ranking the participants in a tournament. *SIAM J. Appl. Math.* 38, 1 (1980), 108–111.
- [21] Joseph Singleton and Richard Booth. 2020. An Axiomatic Approach to Truth Discovery. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS '20)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2011–2013.
- [22] Joseph Singleton and Richard Booth. 2021. Rankings for Bipartite Tournaments via Chain Editing. (2021). arXiv:cs.MA/2101.02476
- [23] Giora Slutzki and Oscar Volij. 2005. Ranking participants in generalized tournaments. *International Journal of Game Theory* 33, 2 (2005), 255–270.
- [24] Giora Slutzki and Oscar Volij. 2006. Scoring of web pages and tournaments—axiomatizations. *Social Choice and Welfare* 26, 1 (2006), 75–92.
- [25] Dong Wang, Lance Kaplan, Hieu Le, and Tarek Abdelzaher. 2012. On Truth Discovery in Social Sensing: A Maximum Likelihood Estimation Approach. In *Proceedings of the 11th International Conference on Information Processing in Sensor Networks (IPSN '12)*. ACM, New York, 233–244.
- [26] Mihalis Yannakakis. 1981. Computing the minimum fill-in is NP-complete. *SIAM Journal on Algebraic Discrete Methods* 2, 1 (1981), 77–79.