


Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>

Contents lists available at ScienceDirect

Artificial Intelligence

www.elsevier.com/locate/artintDesigning competitions between teams of individuals[☆]Pingzhong Tang^{a,*}, Yoav Shoham^{b,c}, Fangzhen Lin^a^a Department of Computer Science, Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong^b Computer Science Department, Stanford University, CA 94305, United States^c Microsoft Israel R&D Center, Herzliya, Israel

ARTICLE INFO

Article history:

Received 15 June 2009

Received in revised form 26 April 2010

Accepted 26 April 2010

Available online 29 April 2010

Keywords:

Team competition

Mechanism design

Truthfulness

Moral hazard

Dominant strategy implementation

ABSTRACT

We consider a setting with two teams, each with a number of players. There is an ordering of all players that determines outcome of matches between any two players from the opposing teams. Neither the teams nor the competition designer know this ordering, but each team knows the derived ordering of strengths among its own players. Each team announces an ordering of its players, and the competition designer schedules matches according to the announced orderings. This setting in general allows for two types of manipulations by a team: Misreporting the strength ordering (lack of truthfulness), and deliberately losing a match (moral hazard). We prove necessary and sufficient conditions for a set of competition rules to have the properties that truthful reporting are dominant strategies and maximum effort in matches are Nash equilibrium strategies, and certain fairness conditions are met. Extensions of the original setting are discussed.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

Once upon a time in ancient China, the emperor Qi threw down the gauntlet to his minister Tian for a horse race. The rule was that each of them would announce a ranking of his three horses and each time the two horses with the same rankings would race. As the story goes, Tian learnt that his best horse was not as good as the emperor's best but better than his second best one, and his second best one was not as good as the emperor's second best but better than the emperor's worst one. Knowing that the emperor would be confident enough to announce the true ordering, the clever minister put forward his worst horse first and his best horse second followed by his second best. As a result, while Tian's worst horse lost badly to the emperor's best horse in the first match, he won the second and third matches nevertheless by taking the advantage of mismatches. Tian explained afterwards his strategy to the emperor and its potential application in military matters and as a result, he was promoted to be the general in chief.

Similar examples abound.¹ A somewhat more recent example is the international team competition of table tennis. The schedule is a modification of the horse racing one by adding two matches between the first player and the second player from each team. Smart coaches can also benefit from strategically reporting the orderings. We will return to these examples later after we formally define the problem, which is henceforth called *team competition problem*.

Competition among teams, each consisting of several players, presents at least two types of challenge. The first regards the desirable outcomes. Typically, the basic information is the relative strength of pairs of players, one from each team.

[☆] An earlier version entitled "Team Competition" appeared in *Proceedings of AAMAS'09*.

* Corresponding author.

E-mail addresses: kenshin@cse.ust.hk (P. Tang), shoham@stanford.edu (Y. Shoham), flin@cse.ust.hk (F. Lin).

¹ For a related example in game theory, the Colonel Blotto game, cf. Section 7.

But how does this information get aggregated to decide the relative merit of the two teams? For example, should a team with one strong player and the rest very weak players beat a team consisting entirely of mediocre players? This amounts to defining the appropriate *social-choice functions* in this domain. One contribution of this paper is to establish criteria for such functions; specifically, we adapted the notions of player anonymity, team anonymity, monotonicity, and Pareto efficiency to this setting.

The second challenge is that relative strengths among players are typically not common knowledge. Each team has private information about its players, and the only objective way of getting this information is to play a match and observe the outcome. Playing all pairwise matches is usually not feasible, and so typically a competition among teams proceeds as follows: The teams announce a ranking of their players, and the organizer schedules individual matches based on these rankings according to a formula announced in advance. The matches then take place, and each match adds a certain score to the team of the winner. The team with the highest aggregate score wins the competition. But this opens the door to two ways in which teams can manipulate the outcome: It can misreport the true ordering of their players, and it can throw a match (that is, deliberately lose). This is the problem of implementing the social-choice function, or of *mechanism design* (see [10,14] for introductions).

Another contribution of this paper is the identification of necessary and sufficient conditions for implementing in *dominant strategies* social choice functions which satisfy the specified axioms. That is, identifying conditions under which it is best for a given team to truthfully reveal the ordering among its players – no matter what the other team does, as well as the conditions under which it is best for a given team to play its best in each match – knowing the other team also plays its best. These results are extended to a more general setting where the outcome of a match between two players is probabilistic – decided according to a *winning probability matrix* [6].

The remainder of the paper is organized as follows: We next formulate the team competition problem as a mechanism design problem, identify the basic forms of mechanisms and state the desirable properties in this domain. In Sections 4 and 5, we characterize the conditions under which the mechanisms satisfy these properties. We then generalize our results in two directions in Section 6 and discuss related work in Section 7. Finally in Section 8, we briefly discuss future research topics related to team competition.

2. Basic models

We now give the mathematic models for analyzing team competition.

2.1. Team competition environments

Team competition environment is the setting where the designer operates.

Definition 2.1. A team competition environment C is a tuple (A, B, Θ, O, R) , where

- $A = \{a_1, \dots, a_n\}$ is the set of players of team A .
- $B = \{b_1, \dots, b_n\}$ is the set of players of team B .
- Θ is the set of possible states, where:
 - Each state $\theta \in \Theta$ uniquely defines a linear order $>_\theta$ on $A \cup B$. If $a >_\theta b$, then a beats b in state θ .
 - We denote by θ_A and θ_B the orderings on A and B that are derived from θ respectively. θ_A and θ_B can be seen as the private information of A and B . We denote by Θ_A and Θ_B the sets of all possible θ_A and θ_B .
- $O = \{(s_A, s_B) \mid s_A, s_B \in \mathbb{R}\}$ is the set of outcomes of the competition. s_A and s_B are the scores for teams A and B , respectively.
- R is a preference relation over O .

We consider R to be the one that team A weakly prefers (s_A, s_B) to (s'_A, s'_B) iff

$$s_A \geq s'_A \quad \text{and} \quad s_B \leq s'_B,$$

and team A strictly prefers (s_A, s_B) to (s'_A, s'_B) iff

$$(s_A > s'_A \text{ and } s_B \leq s'_B) \quad \text{or} \quad (s_A \geq s'_A \text{ and } s_B < s'_B).$$

Team B has the opposite preference. We note that when $s_A > s'_A$ and $s_B > s'_B$, the preference between (s_A, s_B) and (s'_A, s'_B) is not defined.²

An easy way to complete the preference defined above is to restrict a mechanism on certain set of outcomes satisfying, for each state θ , $s_A + s_B = c$, for some constant c . Such a mechanism is called a constant-sum mechanism.

² We will get back to another type of preference where each team is not so sensitive about scores but only cares about winning or losing. In this case, 2 : 3 is as desirable as 3 : 4.

2.2. Strategies

Typically, each team communicates with a mechanism by sending it a message. In team competition context, such a message is confined to the form of an ordering of players of that team. A *strategy* of each team then specifies how to choose among its orderings of players, given its true ordering of strengths.

Definition 2.2. $S_A : \Theta_A \rightarrow L_A$, is the set of A 's pure strategies that map A 's private information to a linear order on A , where L_A is the set of all linear orders on A . Similarly for S_B .

When there is no restriction on Θ , both Θ_A and L_A denote the set of all permutations on A . We use different notations here to clarify that Θ_A is the set of private information (types) based on which A chooses an ordering in L_A to report.

Similarly, we can define the set of A 's *mixed strategies* to be $\sigma_A : \Theta_A \rightarrow \Omega(L_A)$, where $\Omega(L_A)$ is the set of probability distributions over L_A . We assume both teams are risk neutral. As a result, when they play a mixed strategy, the outcome is equivalent to the expected score profile.

2.3. Generalized round-robin mechanisms

Definition 2.3. Given a team competition environment and a message profile (L_A, L_B) reported by A and B , a generalized round-robin mechanism specifies an outcome via a matrix C , where:

- Each entry $c_{i,j}$ in C denotes the score assigned to the match between a_i and b_j , where a_i is the i -th player of L_A and b_j the j -th player of L_B .
- The winner of the match gets $c_{i,j}$ and the loser gets 0.
- The total score that team A can get in state θ is $s_A = \sum_{(a_i >_\theta b_j) \in A \times B} c_{i,j}$. Similar for s_B .
- Such a pair (s_A, s_B) creates an outcome in O .

In comparison with standard mechanism definition [11, Chapter 10], the matrix C plays the role of an outcome function: for each state, the matrix maps the reports from A and B to an outcome $(s_A, s_B) \in O$. Note that $s_A + s_B = c$, where $c = \sum_{1 \leq i, j \leq n} c_{i,j}$, which implies that every generalized round-robin mechanism is constant-sum and the preference relation P over O is complete.

Note that there are potentially n^2 matches since there are n^2 entries in the matrix. However, if $c_{i,j} = 0$, then a match between a_i and b_j is not necessary. Note also that, if $c_{i,j} \neq 0$, it does not necessarily mean that there is only one match between a_i and b_j . It means that the sum of scores of the matches between a_i and b_j is $c_{i,j}$.

Both examples mentioned at the beginning of the paper can be categorized as generalized round-robin mechanisms, with the following score matrices:

Example 2.1 (Horse race).

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

Example 2.2 (Table tennis).

$$\begin{bmatrix} 1 & 1 & 0 \\ 1 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

2.4. Sequential mechanisms

Consider a generalized round-robin mechanism where there are 4 matches: $(a_1$ vs $b_1)$, $(b_1$ vs $a_2)$, $(a_2$ vs $b_2)$ and $(a_1$ vs $b_2)$. Suppose these matches take place sequentially and the first player always beats the second. We can predict after the first three matches that $a_1 >_\theta b_2$, without observing the outcome of the fourth one.

The above intuition can be realized in designing more compact mechanisms called *sequential mechanisms*, where matches take place sequentially and the current match is jointly scheduled by the reported orderings as well as the results of previous matches.

Definition 2.4. Given a team competition environment and a message profile (L_A, L_B) reported by A and B , a sequential mechanism is a tuple (H, f_n, R_s) where

- $H = H_T \cup H_N$ is a set of histories, where H_T denotes the set of terminal histories and H_N the set of nonterminal ones. They are defined inductively as follows:
 - $\emptyset \in H_N$.
 - If $h \in H_N$, then $h :: (a_i > b_j) \in H$ and $h :: (b_j > a_i) \in H$. It says if h is a nonterminal history, by concatenating it with the match where a_i beats b_j or b_j beats a_i , a new history is generated. The new history can be either terminal or nonterminal.
- $f_n : H_N \times L_A \times L_B \rightarrow A \times B$, is a next function that maps each nonterminal history as well as the reported messages to a pair of players to compete in the next match.
- $R_s : H_T \rightarrow O$, is a scoring rule that maps each terminal history to an outcome, that is, a score profile.

Sequential mechanism generalizes round-robin mechanism in a trivial sense that every round-robin mechanism can be represented by a sequential one: the one that schedules a list of independent matches sequentially. Quite often, it is more interesting to focus on certain specific classes of sequential mechanisms. For instance, the following “knock-out” competition is popular in the Go community.

Example 2.3 (*Knock-out competition*). Upon receiving the reported lists $\{a_1, \dots, a_n\}$ and $\{b_1, \dots, b_n\}$,

- a_1 vs b_1 will be initiated as the first match.
- In the following rounds, if the current match is a_i vs b_j , then the next function assigns a_{i+1} vs b_j if b_j beats a_i and a_i vs b_{j+1} otherwise.
- The set of terminal histories are those with every player in one team has lost. The scoring rule assign the winning team n points and the losing team the number of matches that it wins.

There are at most $2n - 1$ matches in a knock-out competition because each match eliminates one player. This property perfectly fits into the context of Go competitions, where a match normally takes hours.

3. Desirable properties

There are several desirable properties a designer wishes to implement. In the horse race story, one might hope for a mechanism that forces both teams to report the truth. Assuming the truth, one might hope that the outcome of the match should fairly reflect the true strengths: it seems fair to anticipate the emperor to win. One might also hope the total number of matches to be as small as possible. We present some of the standards using *social-choice function* and ask what mechanisms truthfully implement the social choice function.

3.1. Social-choice functions

As already mentioned, a social-choice function describes which outcome should occur for a given state.

Definition 3.1. A social choice function $f : \Theta \rightarrow O$ maps a state to an outcome.

3.1.1. Axioms

The first axiom on a choice function f is the so-called player anonymity, which says players are indistinguishable inside a team.

Definition 3.2 (*Player anonymity*). If $p : A \rightarrow A$ is a permutation function and $f(\theta) = o$, then $f(\theta') = o$ where θ' is obtained from θ by replacing each $a \in A$ by $p(a)$. Similar for team B .

For example, state

$$a >_{\theta_1} b >_{\theta_1} b' >_{\theta_1} a',$$

should lead to the same outcome as state

$$a' >_{\theta_2} b >_{\theta_2} b' >_{\theta_2} a,$$

as well as state

$$a >_{\theta_3} b' >_{\theta_3} b >_{\theta_3} a'.$$

The second axiom, team anonymity, says that a choice function f does not discriminate for or against one particular team.

Definition 3.3 (*Team anonymity*). If $p : A \rightarrow B$ is a bijection between A and B and $f(\theta) = (s_A, s_B)$, then $f(\theta') = (s_B, s_A)$ where θ' is obtained from θ by swapping each $a \in A$ and $p(a) \in B$.

For example, again if $f(\theta_1) = (s_A, s_B)$ and θ_1 is as follows:

$$a >_{\theta_1} b >_{\theta_1} b' >_{\theta_1} a',$$

then $f(\theta_4) = (s_B, s_A)$, where θ_4 is as follows:

$$b >_{\theta_4} a >_{\theta_4} a' >_{\theta_4} b'.$$

The third axiom is the so-called monotonicity, which says no worse outcome is brought about for a team if none of its players falls in the overall ranking.

Definition 3.4 (*Monotonicity*). For any two states θ and θ' , if $f(\theta) = o$, $f(\theta') = o'$ and θ' is an improvement to θ for team A , then o' is at least as good as o to A . A state θ' is an improvement to another state θ for team A , if $\forall a \in A$, the ranking of a in state θ' is improved or stays the same as in state θ . Similar for team B .

Finally, the last axiom, called Pareto efficiency, says that if one team has the better i -th best player for all i , then it should get a no lower score in the final outcome.

Definition 3.5 (*Pareto efficiency*). If in any state θ satisfying $\forall 0 \leq i \leq n$, the i -th ranked player of team A is better ranked than team B , then $f(\theta) = (s_A, s_B)$ satisfies $s_A \geq s_B$.

For example, if

$$a >_{\theta_5} b >_{\theta_5} a' >_{\theta_5} b',$$

then $s_A \geq s_B$. Pareto efficiency is not independent of team anonymity and monotonicity since we can prove the latter two properties imply Pareto efficiency.

Proposition 3.1. A social choice function f satisfies Pareto efficiency if it is both anonymous and monotonic.

Proof. Suppose otherwise, then there exists a state θ such that $\forall 0 \leq i \leq n$, the i -th ranked player a_i of team A is better ranked than that b_i of team B and $s_A < s_B$. Now we swap the role of a_i and b_i for all i in θ and we call the new state θ' . By team anonymity, we have $f(\theta') = f(s_B, s_A)$, which is a worse outcome than $f(\theta)$ for B . However, since θ' is an improvement to θ for team B , the new outcome should be no worse for team B by monotonicity. This leads to a contradiction. \square

3.1.2. Examples of social-choice functions

It is not difficult to see that the following four social-choice functions satisfy all the axioms mentioned above.

- Borda Count: Suppose $>_{\theta}$ on $A \cup B$ with $|A \cup B| = 2n$, we assign the top-ranked player $2n - 1$ points, the second-ranked $2n - 2$ points, ..., and the last-ranked 0 point. Let s_A be the sum of points of all the players in A . s_B can be defined symmetrically. $f_{BC}(\theta) = (s_A, s_B)$.
- Horse Race: Suppose in θ , $(a'_1 > a'_2 > \dots > a'_n)$ and $(b'_1 > b'_2 > \dots > b'_n)$. Define $s_A = |\{(a'_i, b'_i) \mid a'_i >_{\theta} b'_i\}|$ and $s_B = n - s_A$. $f_{HR}(\theta) = (s_A, s_B)$.
- Max: Suppose the best players of A and B by θ are a and b respectively, then $f_{Max}(\theta) = (1, 0)$ if $a >_{\theta} b$ and $f_{Max}(\theta) = (0, 1)$ otherwise.
- Min: Suppose the worst players of A and B by θ are a and b respectively, then $f_{Min}(\theta) = (1, 0)$ if $a >_{\theta} b$, $f_{Min}(\theta) = (0, 1)$ otherwise.

In other words, f_{BC} sums up the rankings of all players in each team and f_{HR} the winnings of pairwise comparisons between players of the same rank. In addition, f_{Max} and f_{Min} compare the best and worst players respectively. f_{BC} is a natural extension of the seminal rank-order (aka Borda count) voting rule to the context where voters' objectives are to select a cabinet of n candidates. f_{HR} captures a class of barter auctions where each bidder places n non-monetary, indivisible bids to compete for n equally valuable objects. One fair allocation for such an auction is that of f_{HR} . f_{Max} is widely employed in Olympic competitions such as the Long Jump, when each athlete is thought of as a team, each jump attempt is thought of as an individual player and the gold medal is for the athlete with the longest jump attempt, that is, the choice of f_{Max} . f_{Min} , on the other hand, embodies the preservative attitude of *maxmin*, according to which a set of players are evaluated at their worst player.

3.2. Truthfulness and truthful implementation

Definition 3.6. A mechanism is dominant-strategy truthful if for every state and each team, reporting truthful order yields a no worse outcome than any other order, no matter what the other team does.

An outcome o_1 is worse than o_2 to team A if A strictly prefers o_2 to o_1 . Aware of its private information, each team in such a mechanism would choose to report its truthful ordering because it is in its best interest to do so.

Definition 3.7. A mechanism M truthfully implements a social-choice function f in dominant strategies, if M is dominant strategy truthful and if both teams report truthfully, the resulting outcome coincides with the one prescribed by f .

If a mechanism truthfully implements a choice function in dominant strategies, both teams prefer to report truthfully. Moreover, truthful reports lead to the desirable outcome prescribed by f . For example, f_{HR} and f_{Min} , when thought of as mechanisms, are not truthful while f_{Max} is truthful as a mechanism.

3.3. Frugality

As we shall show, the canonical generalized round-robin mechanism whose score matrix consisted of all 1's, is always truthful but wasteful in terms of number of matches. It is thus reasonable to pursue frugality, that is, when implementing a social choice function, select a mechanism with the smallest number of matches.

4. The results

This section presents answer to the question asked earlier: what mechanisms are truthful and implement the desirable social choice functions.

4.1. Implementation by generalized round-robin mechanisms

We say that a matrix $C_{n \times n}$ is *double-decreasing* if $c_{i_1, j_1} \geq c_{i_2, j_2}$ whenever $i_1 \leq i_2$ and $j_1 \leq j_2$ hold simultaneously.

Theorem 1.

1. A generalized round-robin mechanism M is dominant strategy truthful iff its score matrix $C_{n \times n}$ is double-decreasing.
2. If a generalized round-robin mechanism M truthfully implements a social choice function f in dominant strategies, then
 - f is player anonymous;
 - f is team anonymous iff the score matrix satisfies $C = C^T$, where C^T is the transposition of C ;
 - f is monotonic iff the score matrix of M has no negative entry.

Proof.

1. \Rightarrow : If M is dominant strategy truthful, without loss of generality, suppose there exist i, j such that $c_{i, j} < c_{i+1, j}$. Now consider such a state θ : $b_1 > b_2 > \dots > b_{n-1} > a_1 > \dots > a_i, b_n, \dots, a_n$. In other words, θ is a state where team A can win only i matches against the worst player b_n of B . Now if B reports b_n as its j -th player, then if A reports honestly, he will get $c_{1, j} + \dots + c_{i, j}$ while if A swaps a_i and a_{i+1} , A will get a better score $a_{1, j} + \dots + c_{i-1, j} + c_{i+1, j}$, which contradicts the dominant strategy truthfulness of M .
 \Leftarrow : If $C_{n \times n}$ is double-decreasing, for any state θ and any $b \in B$ reported as j -th player, suppose according to θ , we have $a_1 > \dots > a_i > b > a_{i+1} > \dots > a_n$. If A reports honestly, it will get $c_{1, j} + \dots + c_{i, j}$ from b , otherwise, it will get $c_{m_1, j} + \dots + c_{m_i, j}$. Since C is double-decreasing, we have $c_{1, j}, \dots, c_{i, j}$ are the greatest i entries in column j of C , so $c_{1, j} + \dots + c_{i, j} \geq c_{m_1, j} + \dots + c_{m_i, j}$. Since j and θ are arbitrarily chosen, we have that A is dominant strategy truthful.
2. This part follows from definitions. \square

Note that, the score matrix in neither the horse racing example nor the table tennis example is double-decreasing. According to Theorem 1, they are not dominant strategy truthful. In the table tennis example, if the state is as follows: $a_1 > b_1 > b_2 > a_2 > b_3 > a_3$. Note that if both A and B report truthfully, B would lose the competition with 3 : 2. However, if B misreport his order as $b_1 > b_3 > b_2$ and A still reports truthfully, B would win the competition with 2 : 3.

We also remark that Theorem 1 still holds if we change the solution concept to *ex post* equilibrium.³ For the proof, we only need to adjust the state θ in the proof of Theorem 1 to state θ' as

$$b_1 > b_2 > \dots > b_{j-1} > a_1 > \dots > a_i, b_j, \dots, a_n, b_{j+1}, \dots, b_n.$$

³ *Ex post* truthfulness says that, regardless of the state, it is in each team's best interest to report truthfully, as long as the other team also reports truthfully.

Theorem 2.

1. The generalized round-robin mechanism truthfully implements f_{BC} in dominant strategies if its score matrix $C_{n \times n} = 1_{n \times n}$, where $1_{n \times n}$ is the matrix with every entry being 1⁴;
2. The generalized round-robin mechanism truthfully implements f_{Max} in dominant strategies if its score matrix satisfies $c_{1,1} = 1$ and $c_{i,j} = 0$ otherwise;
3. There is no generalized round-robin mechanism truthfully implements either f_{HR} or f_{Min} in dominant strategies.

Proof.

1. It is not hard to see that for $1_{n \times n}$, the mechanism simply counts the sum of the number of opponents that are weaker for each player. Moreover, it is non-decreasing. So it truthfully implements f_{BC} minus a constant $\frac{n(n-1)}{2}$ in dominant strategies. The constant stands for the sum of additional scores if they are allowed to play with their own team mates.
2. This part follows directly from definition.
3. • Suppose M with score matrix C truthfully implements f_{HR} . Now consider $\theta: a_1 > b_1 > \dots$ where $f_{HR}(\theta) = (s_A, s_B)$, then by only swapping a_1 and b_1 in θ , we obtain $\theta': b_1 > a_1 > \dots$, where $f_{HR}(\theta') = (s_A - 1, s_B + 1)$. This is possible only if $c_{1,1} = 1$. Similarly, consider $\theta'': \dots > a_n > b_n$ where we have $f_{PC}(\theta'') = (s'_A, s'_B)$. Again, by swapping a_n and b_n , we obtain $\theta''': \dots > b_n > a_n$ where $f_{PC}(\theta''') = (s'_A - 1, s'_B + 1)$. This is possible only if $c_{n,n} = 1$. Since M is dominant truthful, we have C is non-increasing, which further implies $C = 1_{n \times n}$. However, $1_{n \times n}$ obviously does not implement f_{HR} . A contradiction.
 • Suppose M with score matrix C truthfully implements f_{Min} . Now consider $\theta: a_1 > b_1 > \dots$ where $f_{HR}(\theta) = (s_A, s_B)$, then by swapping a_1 and b_1 in θ , we obtain $\theta': b_1 > a_1 > \dots$, where $f_{HR}(\theta') = (s_A, s_B)$, as long as $n \geq 2$. This is possible only if $c_{1,1} = 0$. Similarly, consider $\theta'': \dots > a_n > b_n$ where we have $f_{PC}(\theta'') = (s'_A, s'_B)$. Again, by swapping a_n and b_n , we obtain $\theta''': \dots > b_n > a_n$ where $f_{PC}(\theta''') = (s'_A - 1, s'_B + 1)$. This is possible only if $c_{n,n} = 1$. We now have $c_{n,n} = 1 > c_{1,1}$, contradicting to the truthfulness condition. \square

We remark that the current international team competition for tennis (aka Davis Cup) uses a generalized round-robin mechanism whose score matrix is $1_{2 \times 2}$, thus it implements Borda Count.

4.2. Implementation by sequential mechanisms

Unlike a generalized round-robin mechanism, the preference P of which is a total order relation on a constant-sum set of outcomes, there are cases where P is not complete for non-constant-sum sequential mechanisms. Thus, it makes sense to focus on certain specific mechanisms. We now introduce a particular class of sequential mechanisms called “knock-in” competition in contrast to the “knock-out” one mentioned in Example 2.3. In a knock-in competition, the loser of the previous match stays to compete with the winner’s successive team mate.

Definition 4.1 (Knock-in competition). Upon receiving the reported lists of players $\{a_1, \dots, a_n\}$ and $\{b_1, \dots, b_n\}$,

- a_1 vs b_1 is initiated as the first match.
- In the following rounds, let the current match be a_i vs b_j , the next function assigns a_{i+1} vs b_j if a_i beats b_j and a_i vs b_{j+1} otherwise.
- The set of terminal histories are those where every player has won in one team, which we call the winning team. The other team is called the losing team.

Note that the losing team does not necessarily lose the competition in terms of score. The observation below follows immediately from definition.

Observation.

1. Every knock-in competition has a number of at most $2n - 1$ matches.
2. By assigning 1 to the winning team and 0 to the losing team in each terminal history, the knock-in mechanism always yields an outcome that coincides with the one prescribed by f_{Min} , and therefore truthfully implements f_{Min} in dominant strategies.

We remark similarly that a knock-out competition truthfully implements f_{Max} , by assigning 1 to the winning team. However, unlike knock-in competition which can implement many other social-choice functions as well, f_{Max} is essentially

⁴ In fact, it implements Borda Count minus a constant $\frac{n(n-1)}{2}$.

the only desirable choice function that knock-out can truthfully implement, since the team with the best player will always put this player on top and he will beat the whole other team in turn.

The flexibility of designing a knock-in competition lies in the choice of scoring rules. The scoring rule should be designed in a way such that, on the one hand, the preference for the set of terminal histories is well defined and, on the other hand, the scores align with the incentives of truth reporting. There are at least two possible classes of scoring rules.

4.2.1. Score by play order

With this type of scoring rule, we assign a constant score c_1 to the winner of the first match, c_2 the second match, and so on. Score and outcome are defined as usual. Note that the number of matches in such a knock-in competition ranges from n to $2n - 1$. In order to preserve the constant-sum property, we assign $2n$ constants as follows:

Definition 4.2. In a score-by-play-order rule, we have a list of $2n$ constants $\{c_1, c_2, \dots, c_{2n}\}$, we assign c_1 to the winner of the first match, c_2 the winner of the second match and so on. When reaching a terminal history after n_0 matches, $n \leq n_0 \leq 2n - 1$, we assign the remaining constants c_{n_0+1}, \dots, c_{2n} to the remaining players (the order does not matter). In the end, each player receives a score.

In this way, we have $s_A + s_B = \sum_{1 \leq i \leq 2n} c_i$, a constant. The following theorem characterizes the dominant strategy truthfulness of knock-in competitions with this type of scoring rules.

Theorem 3. For a knock-in competition M with a score-by-play-order rule,

1. M is dominant strategy truthful iff $\{c_1, c_2, \dots, c_{2n}\}$ is non-increasing.
2. If M truthfully implements a social choice function f in dominant strategies, then f is player anonymous, team anonymous, monotonic.
3. M truthfully implements:
 - f_{BC} with $\{2n - 1, 2n - 2, \dots, 1, 0\}$;
 - f_{Max} with $\{1, 0, \dots, 0\}$;
 in dominant strategies.

Proof.

1. \Rightarrow : If M is dominant strategy truthful and suppose $\{c_1, c_2, \dots, c_{2n}\}$ is not non-increasing. Without loss of generality, let $c_m < c_{m+1}$. Suppose the m -th match is between a_i and b_j . Now consider a state

$$\theta: a_i > b_j > a_{i+1} > b_{j+1}.$$

If A reports truthfully, the competition goes as follows:

- (a) a_i beats b_j , team A gets c_m ;
- (b) a_{i+1} loses to b_j , team B gets c_{m+1} ;
- (c) a_{i+1} beats b_{j+1} , team A gets c_{m+2} .

If A swaps a_i and a_{i+1} , the competitions goes as follows:

- (a) a_{i+1} loses to b_j , team B gets c_m ;
- (b) a_{i+1} beats b_{j+1} , team A gets c_{m+1} ;
- (c) a_i beats b_{j+1} , team A gets c_{m+2} .

Thus by lying, team A gets a better outcome. A contradiction.

\Leftarrow : We first introduce the notations for this part. Note that given a non-increasing score sequence, a knock-in mechanism can be fully characterized at any stage by a tuple (A', B', c) , where A' and B' are the current lists of players that stay in the competition and c is the current winning score. We denote $U_A(A, B, c)$ as the remaining score that team A will get afterwards. Similar for team B . To show the dominant strategy truthfulness, it suffices to show that $U_A(A_T, B, c_1) \geq U_A(A_U, B, c_1)$ for any A_U , where A_T and A_U are the truthful and untruthful lists of team A , respectively. We now prove it by induction on the size of A and B .

- Base case: It is trivial that we have $U_A(A_T, B, c_1) \geq U_A(A_U, B, c_1)$ when $|A| = 1$ and $|B| = 1$.
- Inductive case: Assume that the claim holds when $|A| + |B| < k$. Now consider the case that $|A| + |B| = k$. Suppose that B reports b_1, b_2, \dots, b_t as its play list and the truthful list for A is $a_1 > a_2 > \dots > a_s$. Compare the truthful and untruthful reports of the first player by A .⁵ There are two cases:
 - Case 1. $a_1 < b_1$
 - (a) If A reported truthfully, b_1 beats a_1 , A would then get a score of $U_A(A_T, \{b_2, \dots\}, c_2)$.

⁵ In principle, we should also compare the case where A truthfully reports a_1 as its first player but misreports some players afterwards, however, this case is covered in the inductive assumption.

(b) If A misreported some other player a_i , b_1 beats a_i again, A would then get a score of $U_A(A_U, \{b_2, \dots\}, c_2)$, where A' denotes the misreported list. We know that $U_A(A_U, \{b_2, \dots\}, c_2) \leq U_A(A_T, \{b_2, \dots\}, c_2)$ by inductive assumption.

We conclude that in Case 1, truthful reporting is no worse than misreporting.

– Case 2. $a_1 > b_1$

(a) If A reported truthfully, a_1 beats b_1 , A would then get a total score of $c_1 + U_A(\{a_2, \dots\}, \{b_1, b_2, \dots\}, c_2)$.

(b) If A misreported some other player a_i , there are two subcases:

* Subcase 1. $a_i > b_1$. a_i beats b_1 , A would then get a total score of $c_1 + U_A(A_U \setminus \{a_i\}, \{b_1, b_2, \dots\}, c_2)$. By inductive assumption, this score is no more than $c_1 + U_A(A_T \setminus \{a_i\}, \{b_1, b_2, \dots\}, c_2)$. In the meanwhile, by knowing $a_i > b_1$, we know that $a_j > b_1$ for all $j < i$. Thus, $c_1 + U_A(A_T \setminus \{a_i\}, \{b_1, b_2, \dots\}, c_2) = c_1 + \dots + c_i + U_A(\{a_{i+1}, \dots\}, \{b_1, b_2, \dots\}, c_{i+1})$. This is the exact score when A reported truthfully since it would win the first i matches in a row.

* Subcase 2. $a_i < b_1$. b_1 beats a_i , A would then get a score of $U_A(A_U, \{b_2, \dots\}, c_2)$, which is no more than $U_A(A, \{b_2, \dots\}, c_2)$ by inductive assumption. We will show in a later section that $U_A(A, \{b_2, \dots\}, c_2) \leq c_1 + U(\{a_2, \dots\}, \{b_1, b_2, \dots\}, c_{i+1})$, which is the definition of *moral hazard freeness*.

We conclude in Case 2 that truthful reporting is also no worse than misreporting.

2. • Suppose that M truthfully implements a social choice function f in dominant strategies and f is not player anonymous, which means there exists one state where by swapping two players' rankings of the same team, f yields a different outcome for the new state. However, as we can see, this is impossible since the flows of the knock-in competition are exactly the same for both states.
 - Similarly for team anonymous.
 - Now consider monotonicity, suppose that M truthfully implements a social choice function f in dominant strategies and f is not monotonic, which means when one player improves his overall ranking while the others remain the same, f yields a worse outcome for that team. This amounts to saying that by giving up certain games and pretending to be a weaker player, one can lead to a better outcome for his team. This contradicts to the fact that M is *moral hazard free*, which will be proved later.
3. This part follows from the conclusion of part 1 as well as the definitions of f_{BC} and f_{Max} .

This completes our inductive proof. \square

For example, to truthfully implement f_{BC} , we first let a_1 play against b_1 with the winner getting $2n - 1$ and the loser staying to compete with the next player of the other team, and so on. Since the earlier matches have higher scores and each player would get a score anyway, each team then would like to win as early as possible, so truthful reporting is no worse than misreporting.

Note also that the same result hold if we replace the solution concept by *ex post* equilibrium, the proof of which is identical.

4.2.2. Score by position

With this type of scoring rule, we assign a score to each match similarly to what we did in the generalized round-robin mechanism, except that each match is asymmetric: the score that a_i gets from winning a_i vs b_j may not be the same as b_j gets from winning the same match. This asymmetry undermines the constant-sum property of the mechanism, however, as we will show, it enables a knock-in competition to truthfully implement f_{HR} . Further, to maintain the team anonymity, we require that the score that a_i gets from beating b_j be the same as the score b_i gets from beating a_j . Therefore, only one score matrix is needed to fully specify a scoring rule.

Definition 4.3. In a score-by-position rule with a score matrix $C_{n \times n}$, for any match a_i vs b_j , a_i gets $c_{i,j}$ if he wins and b_j gets $c_{j,i}$ otherwise.

As we mentioned, in general this type of mechanism is not constant-sum thus the preference relation may not be well defined. However, the main purpose to include it is that it can truthfully implement f_{PC} in dominant strategies with certain restrictions on C .

Theorem 4.

- A knock-in competition M with score-by-position rule is dominant strategy truthful if its score matrix satisfies $\forall 1 \leq i, j, i + 1, j + 1 \leq n$:
 1. $c_{i,j} \geq c_{i,j+1} \geq 0$.
 2. $c_{i,j} \geq c_{i+1,j+1} \geq 0$.
- A knock-in competition M with score-by-position rule truthfully implements f_{HR} in dominant strategies with $c_{i,j} = 1$ if $i \geq j$ and $c_{i,j} = 0$ otherwise.

Proof.

- Note that, at any stage, the competition can be completely characterized by the current lists of players (A, B) , thus the remaining score that team A can get at this point can be denoted by $U_A(A, B)$, since the score matrix $C_{n \times n}$ is always the same one thus can be omitted. We can see that the proof of the first claim is identical to that of Theorem 3, when we replace in the proof of Theorem 3 any instance of $U_A(A, B, c_i)$ by $U_A(A, B)$, c_1 by $c_{1,1}, \dots$, and c_i by $c_{i,1}$. Note also that, in competition with score by position rule, it is also necessary to consider the score U_B of team B when comparing score profiles. However, this part is symmetric to that of U_A .
- First, if a mechanism with the score matrix C such that $c_{i,j} = 1$ if $i \geq j$ and $c_{i,j} = 0$ otherwise, then it satisfies:
 1. $c_{i,j} \geq c_{i,j+1} \geq 0$.
 2. $c_{i,j} \geq c_{i+1,j+1} \geq 0$.
 Thus according to the first claim, it is dominant strategy truthful. The implementation of f_{HR} then follows from the fact that each player can get 1 point iff it beats some higher or equally ranked player. \square

In other words, to implement f_{HR} , we assign each player 1 point if he beats some higher ranked opponent while assign each player 0 point if he loses to or wins against some lower ranked opponent. For example, if a_3 beats b_2 , then a_3 gets 1 and b_2 gets 0, otherwise both a_3 and b_2 get 0. This indicates a higher ranked player can never score by competing with a lower ranked player, but he still has to win so that his lower ranked team mates can have a chance to score.

Till now, we have solved the horse racing problem introduced at the beginning of the paper.

4.3. Knock-in competition minimizes the number of matches

We know that in general a generalized round-robin mechanism schedules $O(n^2)$ matches in comparison to $O(n)$ matches in a knock-in competition to truthfully implement a social-choice function. A natural question is if there exists any other mechanism that can do better. We now give a negative answer by proving that knock-in competition is *worst-case optimal*.

Definition 4.4 (*Worst-case optimal*). A mechanism M is worst-case optimal with respect to a social-choice function f if M truthfully implements f and M schedules the minimum number of matches in its worst case.

In other words, suppose K_M is the minimal number of matches that suffices to guarantee the termination of M for any state. M is worst-case optimal means $K_M \leq K_{M'}$ for any M' .

Theorem 5.

- If a knock-in competition M truthfully implements a social-choice function f that is one-to-one, then M is worst-case optimal with respect to f .
- If a knock-in competition M truthfully implements f_{BC} , then M is worst-case optimal with respect to f_{BC} . The same holds for f_{Max} .

Proof.

- Since f is one-to-one, truthfully implementing f implies figuring out the true state given that the two teams reported truthfully. This is the problem of merging two sorted lists using minimum number of comparisons. Now consider the worst case for any knock-in competition, that is, when players from each team appear alternatively in the state,

$$a_1 > b_1 > a_2 > b_2 > \dots > a_n > b_n.$$

Knock-in competition needs $(a_1, b_1), (b_1, a_2), \dots, (a_n, b_n)$, $2n - 1$ matches to figure out this state. However, to tell this state apart, any of these matches has to be made explicit. Suppose not, say we omit (b_1, a_2) , then we are not be able to tell the original state from the following state

$$a_1 > a_2 > b_1 > b_2 > \dots > a_n > b_n.$$

One can similarly verify other matches. This means knock-in minimizes the number of matches in its worst case (not necessarily the worst case of other mechanism), thus we prove its worst-case optimality. In fact, one can view the knock-in competition as the simulation of the 'merge' procedure in the standard merge-sort algorithm, which is known to be worst-case optimal when merging two sorted lists with an equal number of elements.

- Now we prove that knock-in competition is optimal for f_{BC} . Note that, although f_{BC} is not one-to-one, f_{BC} still produces different values for the two states constructed in the previous part of the proof. Thus, to tell the two states apart, the comparison between (b_1, a_2) still has to be made explicit. One can verify this is true for any other comparison listed above. Thus f_{BC} makes no difference to one-to-one in this sense and the proof above follows. For f_{Max} , knock-in implements it in only one match with score by play order rule $\{1, 0, 0, \dots\}$. \square

Given the theorem above, one might ask if there exists such a one-to-one function that knock-in competition truthfully implements. The answer is affirmative. First note that, knock-in mechanism can truthfully implement any decreasing scoring function – a social-choice function in which each player is assigned a score that is a monotonic function of his rank in the state and the associated outcome is computed by the sum of scores of each team – by trivially assigning each match the score of the corresponding rank. The following decreasing scoring function is one-to-one.

Example 4.1 (*Borda Count with taxes*). We consider a variation of Borda Count by imposing a different *tax rate* of each rank. For instance, suppose each team has two players, the best player gets $4 \times (1 - 4\%)$, the second $3 \times (1 - 3\%)$, the third $2 \times (1 - 2\%)$ and the last $1 \times (1 - 1\%)$. One can verify that this is a one-to-one social-choice function which can be trivially implemented by a knock-in competition.

One implicit assumption that we need to address is that the two teams have an equal number of players. Without the assumption, we do not know if the above theorem still holds, since the knock-in procedure is somewhat redundant for identifying the true state in this case. In fact, Knuth ([7], Chapter 5.3.2, Exercise 8) demonstrated a procedure that can merge any sorted list of length 2 with any another of length 8 in 6 comparisons (while knock-in needs 9 in the worst-case). However, it does not negate our theorem since that procedure is not truthful.

Notice that we do not mention f_{HR} or f_{Min} in our theorem. As we shall introduce in the next subsection, we can do even better than knock-in with respect to f_{HR} and f_{Min} by allowing randomization.

4.4. Randomized sequential mechanisms

As mentioned, if we are allowed to *randomize* the next function (that is, non-deterministically select the next match) in sequential mechanism, we can further design truthful mechanisms that are even more compact than knock-in. For instance, the following mechanism truthfully implements f_{HR} in only $n + 1$ matches.

A reports a_1, \dots, a_n and B reports b_1, \dots, b_n , the mechanism schedules the following $n + 1$ matches:

Example 4.2 (*Implementing horse race in $n + 1$ matches*).

1. a_1 vs b_1 , the winner gets 1 and the loser gets 0;
2. a_2 vs b_2 , the winner gets 1 and the loser gets 0;
- ⋮
- n . a_n vs b_n , the winner gets 1 and the loser gets 0;
- $n + 1$. a_i vs b_j , where a_i and b_j are uniformly selected from A and B respectively, the winner and the loser gets 0.

However if we detect one of the following cases,

1. $j < i$, b_i beats a_i and a_i beats b_j ;
2. $j > i$, b_j beats a_i and a_i beats b_i ;

we eject team B , which is an outcome equivalent to giving team $B -\infty$ in payoff.

Similarly, if we detect one of the following cases,

1. $j < i$, a_i beats b_j and b_j beats a_j ;
2. $j > i$, a_j beats b_j and b_j beats a_i ;

we eject team A .

One can see that if some team takes the advantage of misreporting and win certain game by mismatch, there is always a positive probability ($1/n^2$) that it will be detected in the $(n + 1)$ -th match and gets severe punishment. In this way, this mechanism prevents manipulation. Similarly, we can truthfully implement f_{Min} in two matches.

Example 4.3 (*Implementing Min in 2 matches*).

1. a_n vs b_n , the winner gets 1 and the loser gets 0;
 2. a_n vs b_j or a_i vs b_n , where a_i and b_j are uniformly selected from A and B respectively, the winner and the loser get 0.
- The detection and punishment techniques are similar to Example 4.2.

We can see that, if one team benefits from misreporting in a deterministic sequential mechanism, it can be detected with positive probability using an additional match where the players are randomly chosen. In other words, using this

additional match and the punishment technique, we can assume that both teams report truthfully without any condition on the scoring rules.

However, note also that the assumption of *moral hazard freeness* (introduced in the next section) is essential in this example since one team could throw the detecting (or some other) match to get the other team punished.

5. Moral hazard freeness

One assumption on the environment is that once a state θ defines $a >_{\theta} b$, a always beats b . This is the case when players are non-strategic individuals. For instance, each player is a card in some card game.

However, sometimes a team has another level of strategic behavior by letting its players throw certain matches if it leads to an increase in its score. Similar phenomenon is referred to as *moral hazard* in the *principal-agent* model [9].

For example, if we assign a sufficiently large score to the second match in a knock-in competition, each team's best player would rather pretend to be a weaker player in order to compete the second match. This also happens when there are negative entries in the score matrix of a generalized round-robin mechanism.

The best way to tackle this is to ensure "playing one's best" as a *dominant strategy* of a team. In other words, no matter what the other team does, it is always no worse to do one's best. However, we argue that achieving moral-hazard freeness in dominant strategy is sometimes difficult, if possible. To illustrate, consider in generalized round-robin mechanism a strategy that "if we win the match with the minimal score, we will give up the remaining ones". Obviously, the other team has no incentive to play their best since the current match does not outweigh the sum of the remaining matches. As a compromise, we use the following definition, which amounts to moral-hazard freeness in Nash equilibrium.

Definition 5.1 (*Moral-hazard freeness*).

- For a generalized round-robin mechanism C and a message profile (A, B) , let $U_A(A, B, C)$ be the score A gets when both team compete with their full strength. We say C is moral-hazard free if

$$U_A(A, B, C) \leq U_A(A, B, C') + c_{i,j},$$

where C' differs from C only in that $c'_{i,j} = 0$. Similar condition is required for U_B .

- For a sequential mechanism D with current history h , current players (a_i, b_j) and a message profile (A, B) , denote $U_A(A, B, h)$ the score A gets when both team compete with their full strength. We say D is moral-hazard free if

$$U_A(A, B, h :: (b_j > a_i)) \leq U_A(A, B, h) \leq U_A(A, B, h :: (a_i > b_j)).$$

Similar condition is required to for team B .

It follows from a simple backward induction that the above definition amounts to saying that when the other team always plays its best, it is no worse to play our best. That is, playing one's best is a Nash equilibrium.

The following theorem summarizes the result of moral hazard freeness in both generalized round-robin and sequential mechanisms.

Theorem 6.

1. A generalized round-robin mechanism is moral hazard free iff its score matrix C has no negative entry.
2. A knock-in competition with score-by-play-order rule is moral hazard free iff $\{c_1, c_2, \dots, c_{2n}\}$ is non-increasing.
3. A knock-in competition M with score-by-position rule is moral hazard free if its score matrix satisfies $\forall 1 \leq i, j, i + 1, j + 1 \leq n$,
 - (a) $c_{i,j} \geq c_{i,j+1} \geq 0$,
 - (b) $c_{i,j} \geq c_{i+1,j+1} \geq 0$.

Proof. The theorem is a special case of Theorem 7. \square

It is interesting that for knock-in competitions, the conditions that characterize moral hazard freeness coincide with those of dominant strategy truthfulness. As a result, as long as these conditions are satisfied, the resulting mechanism is both truthful and moral hazard free.

It is also not hard to see that both randomized mechanisms introduced in the previous section are not moral hazard free since there is no incentive for them to play their maximal strength in the detecting match, which, on the contrary, provides incentive for them to throw a match and get the other team punished.

6. Generalization of the results

In this section, we generalize the previous results in two different directions.

6.1. Win–lose–tie

We have considered an outcome to be a pair of real numbers representing the scores that each team will receive at the end of the competition. One could argue that in many cases, what really matters is who wins the competition. In our two-team competition setting, this can be done by assuming three possible outcomes 1 (team A won), 0 (tie), and -1 (team A lost) (cf. [16]).

An interesting question then is how this will affect the results so far. First of all, we notice that instead of changing the set of outcomes, the same effect can be achieved by changing the preference relation P into the following ordering: A strictly prefers (s_A, s_B) over (s'_A, s'_B) iff either

$$s_A > s_B \quad \text{and} \quad s'_A \leq s'_B$$

or

$$s_A = s_B \quad \text{and} \quad s'_A < s'_B$$

and is indifferent to (s_A, s_B) and (s'_A, s'_B) iff either

$$s_A > s_B \quad \text{and} \quad s'_A > s'_B$$

or

$$s_A = s_B \quad \text{and} \quad s'_A = s'_B$$

or

$$s_A < s_B \quad \text{and} \quad s'_A < s'_B.$$

It is similar for team B 's preference.

One can reason that, if a condition is sufficient for dominant strategy truthfulness (or implementation of a social choice function), then it is also sufficient here because the weak preference persists (o_1 is weakly preferred to o_2 before implies o_1 is weakly preferred to o_2 now). However, a previously necessary condition may not hold now. For example, one can verify that the generalized round-robin mechanism with the following score matrix $C_{2 \times 2}$:

$$\begin{bmatrix} 9 & 10 \\ 10 & 0 \end{bmatrix}$$

is dominant strategy truthful although it is not double-decreasing.

6.2. Probabilistic match outcomes

One could also argue that in many realistic cases, it is not always the case that the stronger player deterministically beats the weaker player. This is true, especially in sports where one exciting thing is the uncertainty about the match outcome even though the rankings of players are usually common knowledge.

Probabilistic play models abound. Knuth [8] introduced a model for the *knockout tournament* problem. He assumed a linear ordering x_1, x_2, \dots of strengths among players where x_i always beats x_j when $j \geq i + 2$ and x_i beats x_j only at probability p when $j = i + 1$. Graham et al. [5] introduced an alternative model by assuming x_i beats x_j with probability p for any i, j such that $i < j$. Yet, the most popular one is the so-called *monotonic model*⁶ [6,13], widely adopted in the knockout tournament literature, an area that is concerned with reasonably seeding players in a knockout tournament. It is formally defined as follows,

Definition 6.1 (*Monotonic model*). Suppose all the players are linearly ordered according to their strengths $x_1 > x_2 > \dots > x_n$, there is a probability matrix $P_{n \times n}$, whose entry $p_{i,j}$ specifies the chance that x_i beats x_j in a match. Further, $P_{n \times n}$ satisfies the following constraints,

1. $p_{i,j} + p_{j,i} = 1$,
2. $p_{i,j} \geq p_{j,i}$ if $i < j$,
3. $p_{i,j} \geq p_{i,j+1}$.

⁶ We overload the notation of monotonicity, which could also denote an axiom of social-choice function. Each appearance should be clear from the context.

Consider this model in the context of team competition. Since there are $2n$ players in total, a $2n \times 2n$ matrix is needed to fully specify a state. We can condense this matrix into an $n \times n$ one by eliminating the entries describing matches within the same team. Thus, the above definition can be translated in the context of team competition as follows,

Definition 6.2 (State). A state θ specifies a linear ordering $>_\theta$ on $A \cup B$ and a play matrix $P_{n \times n}$, whose entry $p_{i,j}$ specifies the chance that a_i beats b_j . Further, $P_{n \times n}$ satisfies the following constraints,

1. $p_{i,j} \geq 0.5$ if $a_i >_\theta b_j$,
2. $p_{i,j} \leq 0.5$ if $b_j >_\theta a_i$,
3. $p_{i,j} \geq p_{i,k}$ if $b_k >_\theta b_j$,
4. $p_{j,i} \geq p_{k,i}$ if $a_j >_\theta a_k$.

The definitions for other parts of team competition environment as well as two types of mechanisms remain unchanged.

One should be able to conclude that the linearly ordering $>_\theta$ (not necessarily unique, but equivalent for the purpose of this paper) can be derived from the play matrix $P_{n \times n}$ if $P_{n \times n}$ satisfies the above constraints. Therefore, a state can be completely characterized by its play matrix.

As one can see, as long as a play matrix has no zero entry, any team can completely beat the other. In this subsection, we compute their *expected* scores whenever mentioning truthfulness, that is, the expectation of the sum of scores from individual matches.

For constant-sum mechanisms such as generalized round-robin mechanism as well as sequential knock-in competition with rule of scoring by play order, it is apparent that any expected score profile is still constant-sum. Hence, the preference R defined at the beginning of the paper is still complete for the set of expected score profiles.

Furthermore, to define the corresponding properties of social choice function based on the new definition of state, player anonymity and team anonymity remain unchanged. The corresponding new definitions of monotonicity and Pareto efficiency are as follows,

Definition 6.3 (Monotonicity). For any two states θ and θ' , if $f(\theta) = o$, $f(\theta') = o'$ and θ' is an improvement to θ for team A , then o' is no worse than o for A . A state θ' is an improvement to another state θ for team A , if $p_{i,j} \leq p'_{i,j}$ for all possible i, j . The same holds for team B .

Monotonicity says that if in one team, every player weakly increases his chances against all players in the other team, the team should not end up in a worse score.

Definition 6.4 (Pareto efficiency). If in any state θ , whose matrix P satisfies $\forall i, p_{i,i} > 0.5$, then $f(\theta) = (s_A, s_B)$ implies $s_A \geq s_B$, and vice versa.

We still have that Pareto efficiency can be implied by team anonymity and monotonicity.

Proposition 6.1. A choice function f satisfies Pareto efficiency if it is both team anonymous and monotonic.

Proof. Suppose otherwise that P satisfies $\forall i, p_{i,i} > 0.5$ and $f(\theta) = (s_A, s_B)$ satisfies $s_B > s_A$ in state θ . One can imagine that team B also has a play matrix P' such that $p'_{i,j} = 1 - p_{j,i}$ is the probability that b_i beats a_j . We now transit to another state θ' by letting each $p'_{i,j}$ be $p_{i,j}$. Since $p_{i,j} \geq p'_{i,j}$ (because $a_i > b_i, a_j > b_j$ follows from $\forall i, p_{i,i} > 0.5$), such a transition is an improvement for team B . By monotonicity, we still have $s'_B > s'_A$ in θ' . However, by team anonymity, we have $s'_B = s_A < s_B = s'_A$, a contradiction. \square

We are now ready to derive similar results to what we did in the deterministic case.

Theorem 7. In probabilistic settings, we have:

1. A generalized round-robin mechanism M is dominant strategy truthful iff its score matrix $C_{n \times n}$ is double-decreasing.
2. If a generalized round-robin mechanism M truthfully implements a social choice function f in dominant strategies, then:
 - f is player anonymous;
 - f is team anonymous iff the score matrix satisfies $C = C^T$, where C^T is the transposition of C ;
 - f is monotonic iff the score matrix of M has no negative entry.
3. A generalized round-robin mechanism is moral hazard free iff its score matrix C has no negative entry.
4.
 - A knock-in competition with score-by-play-order rule is moral hazard free iff $\{c_1, c_2, \dots, c_{2n}\}$ is non-increasing;
 - A knock-in competition with score-by-position rule is moral hazard free if its score matrix C satisfies $\forall 1 \leq i, j, i + 1, j + 1 \leq n$,
 - (a) $c_{i,j} \geq c_{i,j+1} \geq 0$,
 - (b) $c_{i,j} \geq c_{i+1,j+1} \geq 0$.

5. For a knock-in competition K with a score-by-play-order rule, K is dominant strategy truthful only if $\{c_1, c_2, \dots, c_{2n}\}$ is non-increasing.
6. For a knock-in competition K , if K truthfully implements a social choice function f in dominant strategies, then
 - f is player anonymous, team anonymous, monotonic.

Proof. See Appendix A. \square

Theorem 7 states that for generalized round-robin mechanisms as well as sequential mechanisms (with scoring by play order), the necessary and sufficient conditions for truthfulness as well as moral hazard freeness remain the same in the generalized probabilistic setting.

7. Related work

7.1. The Colonel Blotto game

We now give a related game called *the Colonel Blotto game*, a classic one that has been put forward for almost a century until recently solved by Roberson [12] under some assumption. The game has found its resemblance and applications to multi-object auctions where agents have budgetary limits [3,15]. It is as follows,

Example 7.1. The Colonel Blotto games.

- A two-person constant-sum game.
- Each player has a fixed number of soldiers and is required to partition all his soldiers to a number of battlefields.
- On each Battlefield, the player that allocate the most soldiers wins.
- The payoff of each player is the number of battlefields won.

The game has several variations include multiple players, incomplete information about other players' budgets and different winning prizes for different battlefields (cf. e.g. [1]). For its most original form, it can be considered as a class of team competitions where each team does not know the other team's partition (thus does not know the strength ordering either). As a result, it is strictly more complex than team competition in the sense that it can be thought of as a two-state game where the first stage is for each team to partition the strength and the second is a specific team competition.

7.2. Mechanism design in sports

Our work can also be categorized as a game-theoretic treatment of designing desirable sports rules. Although it is original in the context of team competition, there has been some parallel work on the tournament design problem. Vu et al. [16] provided a computational treatment of seeding knockout tournament. Their main conclusion is that it is computationally hard to maximize the winning probability of a target player by seeding. Thus, it assures that the designer cannot collude with any player when the number of players is considerably large. Altman et al. [2] alternatively provided a social-choice theoretical treatment of manipulations in tournament. Each tournament is treated as a social-choice function that maps a strength state to a winner. They considered several desirable axioms on this setting and provided some possibility and impossibility results.

8. Future work and concluding remarks

We now consider several topics for future research.

We have considered only mechanisms that require both teams to submit their players lists all at once. Sometimes, however, the teams may regret about their original strategies after observing the results of the first few matches. To accommodate this "change of heart", one can consider using games in *extensive form* where the teams can incrementally add their players' lists during a competition. For example, in a knock-in competition where there is no constraint among players' strengths, at any stage of the competition. One could choose a player has a high chance to beat the current opponent but low chances to beat the other remaining opponents.

It is also interesting to consider a more complicated setting where each match takes *multiple* players. Such examples include various card games such as *Bridge*, in which a card can be considered as a player, and a hand of cards teams. Each round can then be considered as a single match that takes multiple cards with the winner chosen and awarded certain score according to the order of the cards. In this setting, a desirable mechanism would not necessarily need to be truthful, but rather it should encourage strategic behaviors.

Instead of minimizing the number of matches, we can consider other criteria such as maximizing total revenue or revenue per match. In this case, we need to define a model that maps strengths to revenue in each match and then optimize the objective accordingly subject to truthfulness and other constraints. One possible way of implementing this is to use the idea of *automated mechanism design* introduced by Conitzer and Sandholm [4].

To sum up, motivated by real world examples, we have formulated the team competition problem in the framework of mechanism design. We have also proposed two typical forms of team competition, identified the desirable properties that these competitions should satisfy and characterized the conditions under which our mechanisms satisfy the properties. We have further generalized our results in two directions and discussed how randomization can help design more compact mechanisms. For future work, it is worth further exploring sequential as well as randomized mechanism. It is also interesting to deploy knock-in competition in real sports competitions as well as multi-agent scenarios.

Acknowledgements

We are grateful to the AIJ editor in charge and anonymous reviewers for their helpful comments. This work was supported in part by HK RGC GRFs 616707 and 616208 and by NSF grant IIS-0205633-001.

Appendix A. Proof of Theorem 7

Lemma A.1 (Rearrangement inequality).

- If $a_1 \geq \dots \geq a_n$ and $b_1 \geq \dots \geq b_n$, then we have $a_1b_1 + \dots + a_nb_n \geq a_1b_{\sigma(1)} + \dots + a_nb_{\sigma(n)} \geq a_1b_n + \dots + a_nb_1$.
- If $a_1 > \dots > a_n$ and $b_1 > \dots > b_n$, then we have $a_1b_1 + \dots + a_nb_n > a_1b_n + \dots + a_nb_1$, when $n \geq 2$

where σ is any permutation of $\{1, \dots, n\}$.

We now prove 1, 2, that is, the conditions for truthfulness and truthful implementation in generalized round-robin mechanism remain the same.

Proof.

1. \Leftarrow : If the score matrix C is double-decreasing, we show that truthful report is a dominant strategy for team A . From the definition, the final score of team A is the sum of the scores that they get by playing with every players in team B . Suppose team B reports b_j as their i -th choice, we now compare the expected scores team A gets from b_j with and without truthful report.

- Truthful report, that is, to report $a_1 > a_2, \dots, a_n$:

$$\sum_{k=1, \dots, n} p_{k,j} c_{k,i}. \tag{1}$$

- Otherwise, to report $a_{\sigma(1)}, \dots, a_{\sigma(n)}$, where σ is any permutation of $\{1, \dots, n\}$:

$$\sum_{k=1, \dots, n} p_{\sigma(k),j} c_{k,i}. \tag{2}$$

Since $p_{1,j} \geq p_{2,j} \geq \dots \geq p_{n,j}$ by the definition of P , and $c_{1,i} \geq c_{2,i} \geq \dots \geq c_{n,i}$ by the non-increasing property of C , (1) \geq (2) then follows from Lemma A.1. Since i, j here are arbitrarily chosen, we have therefore proved this part.

\Rightarrow : If a generalized round-robin mechanism M is dominant strategy truthful, suppose otherwise that $c_{i,j} < c_{i+1,j}$ for some i, j , we compare the total expected score team A gets with and without truthful report, given team B reports their truthful order $b_1 > b_2, \dots, b_n$.

- Truthful report:

$$\sum_{k=1, \dots, n} \sum_{j=1, \dots, n} p_{k,j} c_{k,j}. \tag{3}$$

- Untruthful report $b_1 > \dots > b_{i+1}, b_i, \dots, b_n$. That is, the one obtained from swapping b_{i+1}, b_i in the truthful report.

$$\sum_{k \neq i, i+1} \sum_{j=1, \dots, n} p_{k,j} c_{k,j} + \sum_{j=1, \dots, n} p_{i+1,j} c_{i,j} + \sum_{j=1, \dots, n} p_{i,j} c_{i+1,j}. \tag{4}$$

We now construct a state θ whose play matrix P satisfying $p_{i,l} = p_{i+1,l}, \forall l \neq j$ and $p_{i,j} > p_{i+1,j}$. It's easy to see that it is a well defined play matrix. Now we get that

$$(3) - (4) = p_{i,j} c_{i,j} + p_{i+1,j} c_{i+1,j} - p_{i+1,j} c_{i,j} - p_{i+1,j} c_{i,j}.$$

Since $c_{i,j} < c_{i+1,j}$ and $p_{i,j} > p_{i+1,j}$, again by Lemma A.1, we have in θ that

$$(3) - (4) < 0.$$

This implies that misreporting yields a better expected score for team A , which contradicts our assumption earlier that M is dominant strategy truthful.

2. This part follows from the definition. \square

Next, we prove 3, 4, that is, the conditions for moral hazard freeness in both generalized round-robin and knock-in competitions still hold.

Proof. For 3, that is the condition of moral hazard freeness in generalized round-robin mechanisms follows directly from the definition.

For the first claim of 4, that is, sequential mechanism with rule of score-by-play-order, we first introduce the notation similar to what we did when proving Theorem 3. After several rounds, the remainder of a knock-in competition can be completely characterized by a triple (A_m, B_n, c_i) where the lists of remaining players of teams A and B are A_m with m players, and B_n with n players respectively and the current winning score c_i . We denote $U_A(A_m, B_n, c_i)$ the expected score that team A gets from the remaining matches of the competition.

\Rightarrow : If a knock-in competition K is moral hazard free and suppose otherwise that $c_i < c_{i+1}$ for some i . Suppose when the current score is c_i , the rest of the competition is $(\{a', a'', \dots\}, \{b', b'', \dots\}, c_i)$. We can construct a state whose play matrix says that a' beats b' and b' beats a'' both at probability 1. Now we can see that if team A play their best for the first two matches, they will end up with $c_i + U(\{a'', \dots\}, \{b'', \dots\}, c_{i+2})$, while they will end up with a higher score $c_{i+1} + U(\{a'', \dots\}, \{b'', \dots\}, c_{i+2})$ if they throw the first match between a' and b' . This contradicts to the fact that K is moral hazard free.

\Leftarrow : We prove this part by induction on m and n for competition (A_m, B_n, c_i) . More specifically, we prove the following hold for all m and n if $c_i \geq c_{i+1}$ for all i :

$$U_A(A_m, B_n, c_i) \leq U_A(A_{m-1}, B_n, c_{i+1}) + c_i, \tag{5}$$

$$U_A(A_m, B_n, c_i) \geq U_A(A_m, B_{n-1}, c_{i+1}). \tag{6}$$

In fact (5) and (6) are equivalent to (7),

$$U_A(A_{m-1}, B_n, c_{i+1}) + c_i \geq U_A(A_m, B_{n-1}, c_{i+1}), \tag{7}$$

since

$$U_A(A_m, B_n, c_i) = p(U_A(A_{m-1}, B_n, c_{i+1}) + c_i) + (1 - p)U_A(A_m, B_{n-1}, c_{i+1}).$$

- Base case: One can easily verify that when $c_i \geq c_{i+1}$ for all i , the above inequations hold for $m' = m, n' = 1$ and $m' = 1, n' = n$.
- Inductive case: Assume (7) (hence (5) + (6)) holds for all m', n' such that $m' + n' \leq m + n - 1$, now we show that it holds for $m + n$. We denote the left-hand side of (7) as LHS, and the right-hand side as RHS. By inductive assumption, we have

$$U_A(A_{m-1}, B_{n-1}, c_{i+2}) + c_i \leq LHS.$$

Similarly, we have

$$U_A(A_{m-1}, B_{n-1}, c_{i+2}) + c_{i+1} \geq RHS,$$

since $c_i \geq c_{i+1}$, we have that $LHS \geq RHS$. The inductive step holds as well.

As to the second claim in 4 concerning score-by-position rule, it follows from the same inductive proof as the above one when we replace any instance of $U_A(A, B, c)$ by $U_A(A, B)$, c_i by $c_{m,n}$ and c_{i+1} by $c_{m,(n+1)}$. \square

Note that this also proves Theorem 6.

At last, we prove 5, 6, the truthfulness conditions in knock-in mechanism still hold.

Proof. \Rightarrow The same counter-example in the proof of Theorem 3 also applies here.

For the \Leftarrow part, we are able to prove that this part holds where each team has two players. For general case, we believe this condition still holds, however we do not have a proof at the moment. \square

References

- [1] T. Adamo, A. Matros, A Blotto game with incomplete information, *Economics Letters* 105 (1) (2009) 100–102.
- [2] A. Altman, A.D. Procaccia, M. Tennenholtz, Nonmanipulable selections from a tournament, in: C. Boutilier (Ed.), *IJCAI*, 2009, pp. 27–32.
- [3] J.-P. Benoit, V. Krishna, Multiple-object auctions with budget constrained bidders, *Review of Economic Study* 68 (2001) 155–179.
- [4] V. Conitzer, T. Sandholm, Complexity of mechanism design, in: *Proceedings of the Uncertainty in Artificial Intelligence Conference (UAI)*, Edmonton, Canada, 2002.
- [5] R.L. Graham, D.E. Knuth, O. Patashnik, *Concrete Mathematics: A Foundation for Computer Science*, Addison–Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1994.
- [6] F.K. Hwang, New concepts in seeding knockout tournaments, *Journal of American Mathematics Monthly* 89 (4) (1982) 235–239.
- [7] D.E. Knuth, *The Art of Computer Programming*, vol. 3: Sorting and Searching, Addison–Wesley, 1973.
- [8] D.E. Knuth, Problem 86-2, a random knockout tournament (with solution), *Journal of American Mathematics Monthly* 93 (1987) 127–129.
- [9] J.-J. Laffont, D. Martimort, *The Theory of Incentives: The Principal-Agent Model*, Princeton University Press, 2001.
- [10] E. Maskin, T. Sjoström, *Implementation Theory*, *Handbook of Social Choice and Welfare*, vol. 1, Elsevier, 2002.
- [11] M.J. Osborne, A. Rubinstein, *A Course in Game Theory*, MIT Press, 1994.
- [12] B. Roberson, The Colonel Blotto game, *Economic Theory* 29 (1) (2006) 1–24.
- [13] A.J. Schwenk, What is the correct way to seed a knockout tournament? *Journal of American Mathematics Monthly* 107 (2) (2000) 140–150.
- [14] Y. Shoham, K. Leyton-Brown, *Multiagent Systems: Algorithmic, Game Theoretic and Logical Foundations*, Cambridge Univ. Press, 2009.
- [15] B. Szentes, R.W. Rosenthal, Three-object two-bidder simultaneous auctions: chopsticks and tetrahedra, *Games and Economic Behavior* 44 (1) (2003) 114–133.
- [16] T. Vu, A. Altman, Y. Shoham, On the complexity of schedule control problems for knockout tournaments, in: *AAMAS'09*, 2009, pp. 225–232.