The Pairwise Exchange Tournament

Kruis, J.; Molenaar, D.; Marsman, M.; Maris, Gunter; van der Maas, H.

DOI:
10.13140/RG.2.2.26027.54569

Link to publication

THE PAIRWISE EXCHANGE TOURNAMENT

By Joost Kruis[*,‖], Dylan Molenaar[†,‖] Maarten Marsman[‡,‖]
Gunter Maris[§,**] and Han van der Maas[¶,‖]

University of Amsterdam[‖] *and ACTNext, by ACT*[**]

The objective of this paper is to present a tournament scheme where the winner is determined as a function of pairwise competition between the players. We show that this tournament can be described by a time homogeneous discrete absorbing Markov chain, and derive the appropriate expressions for the probabilities and expectations of the tournament progression.

[*]Supported by NWO (The Dutch organization for scientific research), No. 022.005.0
[†]Supported by NWO Veni grant, No. 451-15-008
[‡]Supported by NWO Veni grant, No. 451-17-017
[§]Supported by NWO, No. CI1-12-S037
[¶]Supported by NWO Creative Industries grant, No. 314-99-107

1. Introduction. Imagine a tournament for a two-player game with k players, where each player starts with a finite amount of tokens. We assume the game is one of skills such that we can assign a rating to each of the k players. A player with a higher rating is more likely to win a game compared to a player with a lower rating. In each round two players are selected proportional to the number of tokens they posses, where players with more tokens have a higher probability to be selected compared to players with less tokens. Each round the selected players play a single game and the winner of this game gets one token from the losing player. As more rounds are played, skilled players will obtain more tokens and are more likely selected to play in rounds compared to less skilled players.

In one version of the tournament, the players continue to play rounds until one player possesses all tokens, this player then wins the tournament. Several alternative setups are possible, such as: 1) The winner is determined to be the first player that possesses a fixed amount of tokens, less that the total number of tokens in the tournament. 2) The winner is determined to be the player that possesses most tokens after playing a fixed number of rounds, or all tokens before this number of rounds is played.

In the remainder of this paper we formalize the tournament scheme and derive the expectation expressions for the tournament winner and number of rounds played as a function of player skill, token allocation, and tournament setup.

2. Pairwise exchange. The skill of a player is expressed as a single real value δ, and $\boldsymbol{\delta} = [\delta_1, \delta_2, \ldots, \delta_k]$ denotes the vector of ratings for the k players in the tournament. The number of tokens a player possesses is expressed in a single integer v, and $\mathbf{v} = [v_1, v_2, \ldots, v_k]$ denotes the vector of how the tokens are allocated among the k players in the tournament, such that the total amount of tokens in the game denoted by n is given by $\sum_i^k v_i = n$.

For each round two players are selected proportional to the number of tokens each player possesses, specifically, the probability that player i and player j are selected to play a round is:

$$2\,\frac{v_i\,v_j}{n(n-1)}\,, \tag{1}$$

when $i = j$, i.e., a player is sampled twice in the same round and plays against themselves, we use $v_j = v_i - 1$ as we sample without replacement, to prevent double spending.

We use the Bradley-Terry-Luce model for pairwise comparisons (Zermelo, 1929; Bradley and Terry, 1952; Luce, 1959) to express the probability that player j wins from player i as a function of the difference in skill of the players:

$$\frac{1}{1+\exp(\delta_i-\delta_j)}\,. \tag{2}$$

We introduce the $k \times k$ stochastic matrix $\mathbf{p}$, where the off-diagonal element p_{ij} contains the probability of players i and j being selected to play in a round of the tournament, and player j winning the game in that round from player i is

$$2\,\frac{v_i\,v_j}{n(n-1)}\,\frac{1}{1+\exp(\delta_i-\delta_j)}\,, \tag{3}$$

and the diagonal element p_{ii} contains the probabilities of the same player being selected twice in one round, and is

$$\frac{v_i(v_i-1)}{n(n-1)}\,. \tag{4}$$

After a game is finished the winning player receives a token from the losing player, and two new players are selected based on the updated allocation of the tokens $\mathbf{v}$. We will write $\mathbf{p}_{\mathbf{v}}$ to make clear that the entries in this stochastic matrix depend on the values in $\mathbf{v}$.

3. Single round transition properties. Let $N = \binom{n+k-1}{k}$ be the number of possible ways n tokens can be allocated among k players. We will write $\mathbf{v}_t = [v_{t1}, v_{t2}, \ldots, v_{tk}]$, to reflect that the allocation of tokens can change during the tournament, where v_{ti} denotes the number of tokens player i possesses under token allocation $\mathbf{v}_t$. Our interest is in the probability of transitioning from one particular allocation of tokens to another allocation of tokens. We will use $\mathbf{v}_a =$ to refer to the 'from' allocation, and $\mathbf{v}_b$ to refer to the 'to' allocation.

Let $\mathbf{P}$ be a square transition matrix of order N, where element P_{ab} contains the probability of transitioning from $\mathbf{v}_a$ to $\mathbf{v}_b$ in one round of the tournament. The number of transitions we can make with non-zero probability from $\mathbf{v}_a$ to $\mathbf{v}_b$ is limited due to the way we have organized the tournament. A particular token allocation can at most transition into $k(k-1)$ different allocations in one round, namely those where the difference between the two allocations is described by the exchange of a single token between two players. One remaining transition with non-zero probability is located on the diagonal entries of matrix $\mathbf{P}$, and describes the event that the same player is selected twice in one round and the token allocation remains unchanged ($\mathbf{v}_a = \mathbf{v}_b$). The probability transitioning from token allocation $\mathbf{v}_a$ to token allocation $\mathbf{v}_b$ in a single round of the tournament is

$$P_{ab} = \begin{cases} \binom{n}{2}^{-1} \dfrac{\prod_{i=1}^{k} \begin{pmatrix} v_{ai} \\ (v_{ai} - v_{bi})^2 \end{pmatrix}}{1+\exp\left(-\sum_{i=1}^{k} \delta_i (v_{bi} - v_{ai})\right)} & \text{if } \sum_{i=1}^{k} (v_{ai} - v_{bi})^2 \text{ is } 2 \\ \binom{n}{2}^{-1} \sum_{i=1}^{k} \binom{v_{ai}}{2} & \text{if } \sum_{i=1}^{k} (v_{ai} - v_{bi})^2 \text{ is } 0 \\ 0 & \text{otherwise} \end{cases} \tag{5}$$

4. Multiple round transition properties. The stochastic matrix $\mathbf{P}$ describes the tournament progression as a discrete-time Markov chain over the possible token allocation among the players, where one unit of time is one round in the tournament. We use r to denote the number of rounds played in the tournament, with $r \in [0, \infty]$. As $\mathbf{P}$ is time-homogeneous, the probability transitioning from token allocation $\mathbf{v}_a$ to token allocation $\mathbf{v}_b$ after playing r rounds in the tournament is $(\mathbf{P}^r)_{ab}$.

The Markov chain with transition matrix $\mathbf{P}$ is an absorbing chain with k absorbing states, i.e., those states where tokens are allocated to one player, and $N - k$ transient states. $\mathbf{P}$ has canonical form

$$\mathbf{P} = \begin{bmatrix} \mathbf{Q} & \mathbf{R} \\ \mathbf{0} & \mathbf{I_k} \end{bmatrix}, \tag{6}$$

where $\mathbf{Q}$ contains the transition probabilities from one transient state to another transient state, $\mathbf{R}$ contains the transition probabilities from the transient states to absorbing states and $\mathbf{I_k}$ is an identity matrix of order k representing that once an absorbing state is entered it cannot be left, and $\mathbf{0}$ is the zero matrix, which represents the transition probabilities from an absorbing state to a transient state.

Rearranging $\mathbf{P}$ in its canonical form allows us to derive the expected progression of the Markov chain more easily, as demonstrated in, for example, Diederich and Busemeyer (2003). The tournament ends when the Markov chain transitions into an absorbing state from a transient state, and the winner is determined by the token allocation of the absorbing state. We use x to denote the player that wins the tournament, where $x \in [1, k]$. Furthermore, we use $\mathbf{z} = [z_1, z_2, \ldots, z_N]$ to denote the probability distribution over the token allocations with which the tournament starts, where z_t denotes the probability of starting the tournament with token allocation $\mathbf{v}_t$. This vector of probabilities $\mathbf{z}$ can be divided into the part containing the probabilities of starting in a non-absorbing token allocation ($\mathbf{z}_q$), and the probabilities of starting in an absorbing token allocation ($\mathbf{z}_i$).

5. Tournament progression; probabilities and expectations. When token allocation at $r = 0$ is non-absorbing, we need to play at least one round before the end of the tournament. The joint probability of player x winning the tournament after playing round r is

$$p(\mathrm{x},\, \mathrm{r} \mid \mathrm{r} > 0) = \mathbf{z}_q\, \mathbf{Q}^{r-1}\, \mathbf{R}_x\,, \tag{7}$$

where $\mathbf{R}_x$ is the column of $\mathbf{R}$ that has the transition probabilities to the absorbing state where all tokens are allocated to player x. Throughout the paper we adopt the practice that $\mathbf{Q}^0 = 1$. The joint probability of player x winning the tournament after playing only one round simplifies then to

$$p(\mathrm{x},\, \mathrm{r} \mid \mathrm{r} = 1) = \mathbf{z}_q\, \mathbf{R}_x\,. \tag{8}$$

The probability that player x wins because the tournament started in the absorbing state where all tokens are allocated to player x is

$$p(\mathrm{x},\, \mathrm{r} \mid \mathrm{r} = 0) = \mathbf{z}_i\, \mathbf{I}_\mathbf{x}\,, \tag{9}$$

where $\mathbf{I}_x$ is the column of $\mathbf{I}_k$ that represents the absorbing state where all tokens are allocated to player x. From these expression we can derive the (marginal) probability that player x wins the tournament as

$$p(\mathrm{x}) = \mathbf{z}_i\, \mathbf{I}_\mathbf{x} + \mathbf{z}_q \sum_{r=0}^{\infty} \mathbf{Q}^r\, \mathbf{R}_x\,, \tag{10}$$

and using geometric series to rewrite the infinite sum over $\mathbf{Q}^r$ we can simplify this expression as

$$p(\mathrm{x}) = \mathbf{z}_i \, \mathbf{I}_{\mathbf{x}} + \mathbf{z}_q \, (\mathbf{I} - \mathbf{Q})^{-1} \, \mathbf{R}_x \, . \tag{11}$$

The (marginal) probability of ending the tournament after round r is

$$p(\mathrm{r}) = \mathbf{z}_q \, \mathbf{Q}^{r-1} \sum_{x=1}^{k} \mathbf{R}_x \, , \tag{12}$$

and the probability of ending the tournament within $u \in [1, \infty]$ rounds is

$$p(\mathrm{r} \leq \mathrm{u}) = \mathbf{z}_i \sum_{x=1}^{k} \mathbf{I}_x \, + \mathbf{z}_q \sum_{r=1}^{u} \mathbf{Q}^{r-1} \sum_{x=1}^{k} \mathbf{R}_x \, , \tag{13}$$

and we write the probability that player x wins the tournament within u rounds as

$$p(\mathrm{x}, \, \mathrm{r} \leq \mathrm{u}) = \mathbf{z}_i \, \mathbf{I}_x \, + \mathbf{z}_q \sum_{r=1}^{u} \mathbf{Q}^{r-1} \, \mathbf{R}_x \, . \tag{14}$$

Assuming that player x wins, the expected duration of the tournament in number of rounds is

$$E(\mathrm{r} \mid \mathrm{x}) = \frac{1}{p(\mathrm{x})} \left[\mathbf{z}_q \sum_{r=1}^{\infty} r \, \mathbf{Q}^{r-1} \, \mathbf{R}_x \right] \, , \tag{15}$$

using again geometric series to rewrite the infinite sum over $r \, \mathbf{Q}^{r-1}$ we can simplify this expression as

$$E(\mathrm{r} \mid \mathrm{x}) = \frac{1}{p(\mathrm{x})} \, \left[\mathbf{z}_q \, (\mathbf{I} - \mathbf{Q})^{-2} \, \mathbf{R}_x \right] \, . \tag{16}$$

The expected duration of the tournament in number of rounds, regardless of the winner, is

$$E(\mathrm{r}) = \sum_{x=1}^{k} \mathbf{z}_q \, (\mathbf{I} - \mathbf{Q})^{-2} \, \mathbf{R}_x \, . \tag{17}$$

6. Tournament progression restrictions. We can limit the expected duration of a tournament by setting a maximum number of rounds to be played, and/or reducing the number of tokens required to be possessed by a single player in order to win the tournament. Let u denote the maximum number of tournament rounds to be played. Then the probability of having no tournament winner after playing u rounds is $1-p(\mathrm{r} \leq \mathrm{u})$. If this situation occurs we can select the player with the highest number of tokens as winner, or sample a winner proportional to the token allocation after playing round u. Let $\mathbf{v}_q$ denote the set of non absorbing token allocations, i.e., the transient states. The probability of being in the transient state corresponding with token allocation $\mathbf{v}_t \in \mathbf{v}_q$ after playing round u in the tournament is

$$p(\mathbf{v}_t \mid \mathrm{r} = \mathrm{u}) = \frac{1}{1 - p(\mathrm{r} \leq \mathrm{u})} \, \mathbf{z}_q \, \mathbf{Q}^u \,, \tag{18}$$

such that the expected number of tokens player x possesses after round u is then

$$E(\mathrm{v}_{rx} \mid \mathrm{r} = \mathrm{u}) = \sum_{\mathbf{v}_t \in \mathbf{v}_q} p(\mathbf{v}_t \mid \mathrm{r} = \mathrm{u}) \, \mathrm{v}_{tx} \,, \tag{19}$$

where v_{rx} denotes the number of tokens allocated to player x after playing round r, and v_{tx} is the number of tokens player x possesses in token allocation $\mathbf{v}_t$. The probability that player x has the highest number of tokens after playing round u, is the sum over $p(\mathbf{v}_t \mid \mathrm{r} = \mathrm{u})$ for all $\mathbf{v}_t$ where $\mathrm{v}_{tx} > \mathrm{v}_{t!x}$. If we opt to sample a winner proportional to the token allocation after playing round u, the probability that player x wins the tournament is

$$p(\mathrm{x} \mid \mathrm{r} = \mathrm{u}) = \frac{1}{n} \, E(\mathrm{v}_{rx} \mid \mathrm{r} = \mathrm{u}) \,. \tag{20}$$

The second option is to reduce the number of tokens a player should posses before winning the tournament. Let $y \in [1, n]$ denote the number of tokens a player should posses to win the tournament. We can establish the new transition matrix $\mathbf{P}^y$, where P^y_{ab} contains the probabilities of transitioning from token allocation $\mathbf{v}_a$ to token allocation $\mathbf{v}_b$ in one round, given that the tournament stops as soon as y tokens are allocated to a player;

$$
P^y_{ab} = \begin{cases}
\binom{n}{2}^{-1} \dfrac{\prod_{i=1}^{k} \begin{pmatrix} v_{ai} \\ (v_{ai}-v_{bi})^2 \end{pmatrix}}{1+\exp\left(-\sum_{i=1}^{k} \delta_i (v_{bi}-v_{ai})\right)} & \text{if } \sum_{i=1}^{k} (v_{ai}-v_{bi})^2 \text{ is 2 and } \sum_{i=1}^{k} \begin{pmatrix} v_{ai} \\ y \end{pmatrix} = 0 \\
\binom{n}{2}^{-1} \sum_{i=1}^{k} \binom{v_{ai}}{2} & \text{if } \sum_{i=1}^{k} (v_{ai}-v_{bi})^2 \text{ is 0 and } \sum_{i=1}^{k} \begin{pmatrix} v_{ai} \\ y \end{pmatrix} = 0 \\
1 & \text{if } \sum_{i=1}^{k} (v_{ai}-v_{bi})^2 \text{ is 0 and } \sum_{i=1}^{k} \begin{pmatrix} v_{ai} \\ y \end{pmatrix} \neq 0 \\
0 & \text{otherwise}
\end{cases} \tag{21}
$$

By rearranging $\mathbf{P}^y$ in its canonical form we can use the expressions from section 5 to derive the probabilities and expectations for the tournament progression, given that the tournament stops as soon as y tokens are allocated to a player. Naturally, we can combine both methods from this section, by applying the expression for a maximum number of rounds to transition matrix $\mathbf{P}^y$.

References.

Bradley, R. A. and Terry, M. E. (1952). Rank analysis of incomplete block designs. I. The method of paired comparisons. *Biometrika* **39** 324–345. MR0070925

Diederich, A. and Busemeyer, J. R. (2003). Simple matrix methods for analyzing diffusion models of choice probability, choice response time, and simple response time. *J. Math. Psych.* **47** 304–322. MR1982565

Luce, R. D. (1959). *Individual choice behavior: A theoretical analysis*. John Wiley & Sons, Inc., New York; Chapman & Hall, Ltd., London. MR0108411

Zermelo, E. (1929). Die Berechnung der Turnier-Ergebnisse als ein Maximumproblem der Wahrscheinlichkeitsrechnung. *Math. Z.* **29** 436–460. MR1545015