The NZASP Rating System

As from 1 January 2022

Overview

The New Zealand Scrabble rating system is similar to most other rating systems, in that it is essentially a modification of the US Chess rating system developed by Arpad Elo and adopted in 1960. This kind of rating system makes use of two basic formulas.

The first formula, known as the ratings curve, defines the proportion of wins expected for each of two players in a given pairing, given the difference between the initial ratings of the two players. For example, if the initial ratings of the two players differ by 100, then using the curve from the NZ system, the higher-rated player would be expected to win around 58% of such games, while the lower-rated player would be expected to win the other 42%. If the two players were to play 14 games against each other, their tournament expectancies would be [0.579 * 14 =] 8.1 out of 14 and [0.421 * 14] = 5.9 out of 14.

The second formula, known as the k-factor, defines the number of points that a player gains or loses per game's difference between their expectancy and their result, based on the player's initial rating. If the two players in the example above were rated 1000 and 900, and each won 7 of the 14 games, the player on 1000 would have a k-factor of 30, so would lose [(7-8.1)*30 = -33] 33 ratings points and have a final rating of 967. The player on 900 would have a k-factor of 31.5, so would gain [(7-5.9)*31.5 = 34.65] 35 ratings points and have a final rating of 935. If the two players were on 2000 and 1900, their k-factors would be 15 and 16.5 respectively. If a player's new rating is less than zero, it is rounded up to zero.

There are no other formulas used in the NZ system, but the two formulas are applied in two different ways, depending on whether the rating being calculated is based on an existing rating, or is being calculated as a new rating, based only on the current tournament. For existing ratings, the formulas above are applied a single time to the player's initial rating. For new ratings, the inverse of the first formula is applied repeatedly (using a scaled win rate to avoid infinities), with the initial rating being replaced by the resulting rating each time, until the resulting rating stops changing. This means that the final new rating is the rating that if the player had started the tournament with that rating, their performance in the tournament would have meant that their provisional rating would not have changed.

Grading

How a player may be placed into a grade for a tournament, and the way in which their rating is calculated as a result, depends on their rating status.

For most players entering a tournament, their rating status is (established) (usually shown as a blank rating status) – these players must be graded strictly by their initial rating (and, for two players on the same rating, by their seeding as shown in the All Ratings list – this will be the order the players were in the last time they had different ratings). These players will be classed as having a ranking if they have played 40 or more games in the previous two years (of weeks in which tournament play was not suspended), but this does not affect the rating process (except where it affects who qualifies for a tournament).

Players who have never played in a tournament rated in the NZ system will have a ratings status of (new). The tournament organiser should place the new player in whichever grade they believe is most appropriate to the player's playing strength, and the player will be rated accordingly.

Players who start a tournament having played at least one rated game previously, but fewer than 30, will be classed as having a provisional rating, i.e. they will have a rating status of (prov).

Players who start a tournament having played at least one rated game previously, but have played no rated games in the previous two years, will be classed as having a historical rating, i.e. they will have a rating status of (hist).

Players may have both (prov) and (hist) rating statuses at the same time (prov/hist). Players with either or both of these statuses will usually be graded according to their current rating, and rated accordingly. But, if the player and the tournament organiser both agree that the player's (prov) or (hist) rating is a sufficiently inaccurate measure of the player's current playing strength to warrant placing the player in a different grade (either higher or lower than that indicated by the player's current rating), then the tournament organiser may do so. Such players will have a rating status of (regraded) for that tournament, and will be rated as if they are a (new) player.

The ability to regrade a (prov) or (hist) player only applies to tournaments where there is more than one grade, and where the players stay in the same grade for the whole tournament. I.e. (prov) or (hist) players can't be regraded for single-division swiss draw, promotion-demotion, or similar tournaments.

Technical details

The sections above cover all the essential elements of the system. This section contains a bit more detail, for those interested.

Ratings curve

From 1999 to 2016 the NZ system used a ratings curve based on the cumulative normal distribution. From 2017, a scaled logistic curve has been used. This is the same ratings curve that is used in the WESPA and North American ratings systems, among others.

The formula that derives a game expectancy from this curve is:

$$1/(1 + EXP((PlayerRating - OpponentRating)/-313))$$

where EXP(x) = e to the power of x.

The inverse of this formula, used for provisional ratings is:

$$-LOG(1/winRate-1)*313$$

where LOG(x) = is the base e logarithm of x.

The win rate is scaled as follows:

$$((wins - games/2) * (games-2) / games + games/2) / games$$

k-factor (multiplier)

The k-factor used in the NZ system from 1999 to 2016 can be restated as:

$$((3000-startRating)/100) * 2$$

In 2017 and 2018, this was modified to:

$$((3000-startRating)/100) * (gamesInTourney/10)$$

Since the beginning of 2019, the k-factor for the NZ system has been:

$$((3000-rating)/100) * 1.5$$

This gives a k-factor of 45 for a zero rating, through to 15 for a 2000 rating.