

## Rating Adjustments

We maintain a ratings list as a means of objectively, and as accurately as possible, reflecting all players' playing standards. This is normally done by recording each match result and increasing the standing of the winning player's rating and reducing that of the losing player. The detailed methods for doing that is described [here](http://www.ttmanawatu.org.nz/Ratings_Explained.pdf) [http://www.ttmanawatu.org.nz/Ratings\\_Explained.pdf](http://www.ttmanawatu.org.nz/Ratings_Explained.pdf). Sometimes we also need some extra adjustments, and they are described here.

Before starting that discussion it needs to be remembered that ratings are completely about reflecting the playing standard of players and nothing else. It is not a reward system. Hence, players should not get a higher or lower rating as a consequence of playing more or fewer games. Or, put differently, let us take two players who play against other players with similar ratings. One player did not play much and has a 1:1 record, while the other player played lots and had a 10:10 record. Based on those results, it seems that the ratings of both players were spot-on and did not need to be adjusted. There is no special treatment for the player who played more matches.

It also means that the points gained by one player match the points lost by another player. That all works fine if there are no players entering or leaving the system and as long as there are enough matches between all players for anyone's advancement or lack of form to be reflected in changed ratings. However, there are some players who play regularly and others who play only occasionally, and some may not have played at all for a year or more. And there are new players coming into the system all the time. They play for some time, mostly improve and some of them then stop playing.

And that causes a problem! And here is why.

Some players get better and advance their rating. Fine! But they can only advance their rating by getting points off the players they play against. That drags down the ratings of those other players even though the playing standard of the other players may not have changed. So, the new players drag everyone else down, something we call 'ratings drift'.

Well, they don't drag down everyone's rating. The new players drag down the rating of the players against whom they played. And those other players with lowered ratings then play someone else and drag them down. So, everyone gets dragged down in the end - everyone, except those who did not play any matches at all. The players who had a rest for a season have no change to their rating. So, the more matches a player plays, the more that player's rating gets dragged down by all the improving new players. And that ratings drift is not fair and needs to be corrected for.

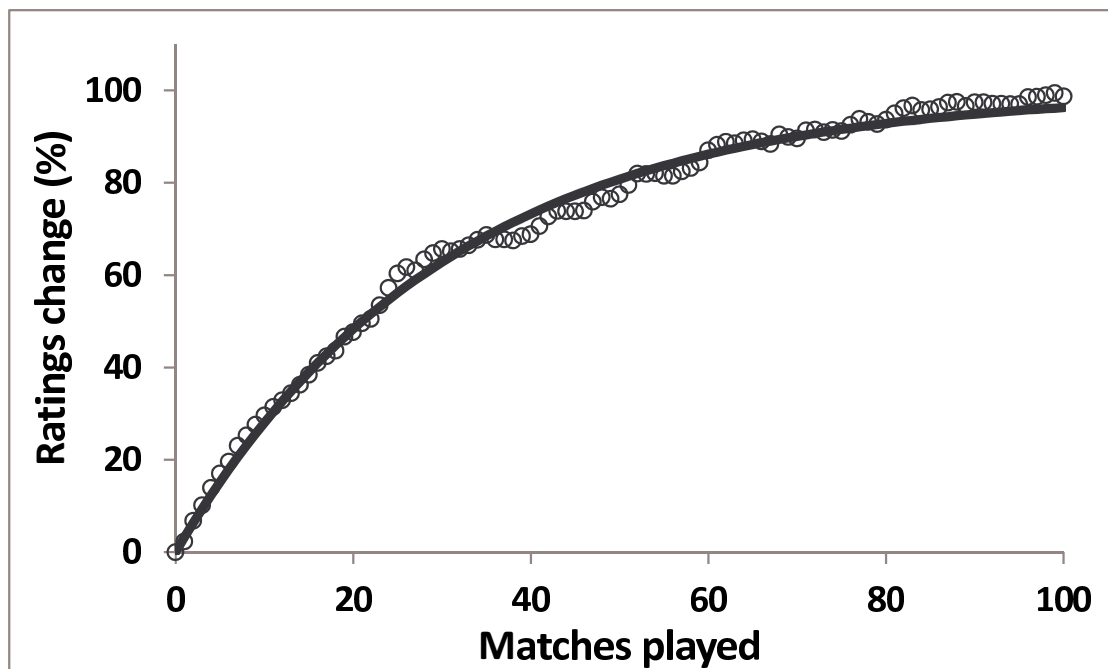
So, the extremes are easy. Someone who has not played at all for a year requires no correction, and someone who has played the whole season needs the biggest correction. But what about the players in-between? How much correction should we apply to those players who have played only one match, or two, or 20? That's where it gets complicated.

Let us remind ourselves what we are trying to achieve here. At the beginning of a season everyone's ratings might have been spot on. Good! During the year, however, some players improve and thereby drag down the ratings of the players they play against. So, we need a compensation for the rating of players to be dragged down simply by other players getting better. So, the question to ask is: by how much does it drag down a player's rating per match played? If we know that we can compensate for it.

So, I have used two approaches: an analytical and a simulation approach. For the analytical approach I have looked at the ratings formula ([refer back to the ratings description for details](http://www.ttmanawatu.org.nz/Ratings_Explained.pdf) [http://www.ttmanawatu.org.nz/Ratings\\_Explained.pdf](http://www.ttmanawatu.org.nz/Ratings_Explained.pdf)) and worked out the probability of certain results occurring and by how much it would shift the calculated ratings. It came out to suggest that ratings shift by about 4% for each match played against players within one's ratings window (opponents with +/- 250 points of a player's own rating).

I also set up a computer program to simulate changes in ratings over time. The average result for 100 trials and for a ratings difference of 50 points is shown in Figure 1. It largely gave the same answer as the analytical solution, but with a slightly lower adjustment constant of only 3.3% per match. It shows how the adjustment is reasonably fast at the beginning (3-4% per match), but as a player's ratings gets closer and closer to the target 100% change, further changes become slower and slower. So, while it makes a lot of difference whether someone has player 5 or 10 matches, it makes little difference whether someone has played 50 or 55 matches.

**Figure 1.** Simulated rating changes per match played. The open symbols are based on 100 simulations with modelled ratings changes. The solid line is a negative exponential relationship given by:  
 $C = 100 (1 - e^{-0.033M})$  where C is the ratings change and M the number of matches played.



Another factor to consider is that the ratings change for which we need an adjustment occurs throughout the season so that matches played at the beginning of the season are still against opponents unaffected by the ratings drift. So, matches played later in the season could be counted more than matches played earlier, but that might start to become a bit too detailed and overly sophisticated.

It is important, however, to consider only matches played against opponents within the ratings window because it makes no difference how many times a player might have played against much stronger or much weaker opponents. It is only through matches within the ratings window that a ratings adjustment is possible, be that a warranted adjustment in line with changing player standards or an unwarranted adjustment as a result of ratings drift.

So, to adjust for ratings drift, the system records the number of matches played by players against opponents within their ratings window. Occasionally, the ratings are then adjusted to bring the average ratings for all players back to close to 1000 points. That adjustment is made for each player based on the formula given in the legend of Figure 1.

It gives a bigger adjustment for players with more matches although for the players who played lots of matches the difference between them becomes pretty small. This greater adjustment for players with more matches is not a reward for having played more matches but a compensation for the ratings drift that would have dragged down the rating of the players who played more matches.