How the TRC Global Rankings work

Technical specifications and detailed overview

1.0 Motivation

TRC Global Rankings provide the sport of horse racing with a global classification of the humans involved in riding, owning and training the best racehorses on the planet. There is also a category of sires.

TRC Global Rankings are different and more powerful than those for other sports because they are designed to be predictive. Like other ranking systems, such as those for golf and tennis, they are an objective, critical assessment of past results intended to reflect in some way the established world order. However, *TRC Global Rankings* are equally forward looking, and they use machine-learning techniques to understand what is important in projecting a competitor's future success.

This last point is the most crucial to appreciate: *TRC Global Rankings* continually test their own ability to predict the next set of Group- and Graded-race results and attempt to minimise the number of so-called ranking violations taking place in future.

A ranking violation occurs when any higher-ranked jockey, owner, trainer or sire is associated with, or responsible for, any horse who finishes behind a lower-ranked rival. Each time this happens, the *TRC Global Rankings* algorithm iteratively makes a small adjustment to its weighting of various factors associated with the individuals concerned – if one is justified.

If, for example, a lower-ranked Canadian horse defeats one from the USA, the *TRC Global Rankings* system might iteratively re-weight the quality of racing in the two countries. But it is crucial not to overfit the data when the results of all races contain a random component, not to mention the fact that the racing environment is dynamic and the fortunes of jockeys, owners, trainers and sires all depend on each other.

So, *TRC Global Rankings* are not some abstract exercise in playing with numbers. They are focussed on a single, simple objective: to find the optimal arrangement of competitors in each category, so that when two ranked individuals meet in competition, we think we know who is more likely to win.

It is this mission statement that drives the underlying mathematics of *TRC Global Rankings*, all of which will be explained underneath. And the consequence of this approach should silence one of the most persistent, and often justifiable, criticisms of rankings: *how on Earth can you compare individuals from different places?*

We can because Australian-trained horses sometimes compete in Hong Kong, and Hong Kong horses sometimes compete against French ones who are sometimes ridden by jockeys who compete in Japan against jockeys who sometimes ride in Britain against horses owned by powerful Middle Eastern connections who are sired by stallions who sometimes have runners in the USA …

You get it. And, if you know anything of the maths that underpin rankings and ratings nowadays, you may recognise a perfect application for techniqueswhich power search engines, recommend movies and even determine the constituents of the perfect wine.

So long as the results database contains sufficient connections between individuals, you really can make the comparisons needed to predict the likely winner of any paired comparison between jockeys, trainers, owners and sires in Group or Graded races around the world – no matter how disparate their location, no matter how different their context. And, there has never been a more interesting time to make those comparisons, with global horse racing's continued march towards internationalisation.

Of course, you may be surprised by, or downright disagree with, the classifications *TRC Global Rankings* produce each week. And that's to the good because, this being horse racing, you don't have to wait very long for your opinion to be tested on the racecourses of the world. And, if it turns out you were right, the rankings will move towards reflecting your opinion, because finding winners is their objective and they learn from their mistakes. The intention, however, is to make the fewest mistakes possible, so *TRC Global Rankings* are very difficult to beat.

2.0 How to interpret the weekly rankings

Each week on a Thursday, a new release of the *TRC Global Rankings* will appear on ***thoroughbredracing.com***. Various drill-down operations are possible.

There are four categories for which a Top 500 is produced, viz

JOCKEYS

OWNERS

TRAINERS

SIRES

Competitors earn ranking **points** via both their past exploits and their projected ability to sustain (or improve on) success. The scale is chosen so that 1000 represents an arbitrary threshold above which the competitor is considered by us within the 'global elite'. This really isn't important, however, but two other aspects related to a competitor's **points** score are:

i) The score determines the within-category ranking, with higher **points** scores corresponding with better (i.e. numerically lower) rankings. For instance, in the first-ever edition of the rankings on January 5, 2014, the #1 ranked individual in the **TRAINERS** category was U.S.-based **Todd Pletcher** with 1050**points**, while the #2 ranked trainer was Ireland's **Aidan O'Brien** with 1037 **points**. In total, there were 12 trainers with a ranking in excess of 1000.

ii) The score also enables a between-categories comparison, for **points** are scaled to reflect the same 'achievement' or ‘dominance’ respective to the population of competitors in the same category. Again, you really can compare.

2.1 How to interpret changes to the weekly rankings

When the best races around the world are run each week, it is inevitable that ranking violations will take place. Nearly every race will see the finishing order violate the predicted order of finish in one or all of the *TRC Global Ranking* categories of **JOCKEYS**, **OWNERS**, **TRAINERS** and **SIRES**.

The rankings algorithm then considers the result of every paired match-up incoming, and cares more about some being wrong than others.

i) It cares more about Group and Grade 1 races than Group 2, and cares more about Group 2 than Group 3;

ii) It cares much more about the precise order of the placed horses than the order of those who trail in towards the rear.

Each week, the rankings algorithm makes changes to **points** as detailed below. It does so to correct its own inevitable errors.

In considering these each week, it is highly important to appreciate this:

BECAUSE THE RANKINGS ARE PREDICTIVE, THEY CONTAIN AN IMPLICIT EXPECTATION OF SUPERIOR PERFORMANCE

If a jockey, trainer, owner or sire is No. 1, the rankings algorithm expects the competitor to win or run close-up in big races every week. *So, a #1 ranked competitor will gain fewer **points** (and sometimes none at all) for the same performance that might cause the #354 ranked competitor to gain 15 **points**.*

Changes to the rankings each week are detailed in the **CHANGE** column on the extreme-right of the display.

All **CHANGE** greater than +4 will result in a competitor appearing in the ***CLIMBERS*** page.

Again, don't be surprised if some of the bigger ***CLIMBERS*** each week are jockeys, owners, trainers and sires you may not be that familiar with.

Neither was the rankings algorithm, which is why it is making the big **CHANGE** !!

3.0 The components that drive the rankings

The **points** score that determines *TRC Global Rankings* is driven by just two components:

i) The *Racing Post Ratings* earned by the competitor's runners, denoted by the column **tRPR**;

ii) The win-loss record of the competitor's runners, denoted by the column **IV**.

Let's look at the two components in more detail:

3.1 How the rankings use *Racing Post Ratings*

tRPR stands for 'translated (or time-decayed) *Racing Post Rating*' (both apply). While *Racing Post Ratings* (RPR) are a widely accepted measure of racehorse merit, they need certain modification to make them fit for our purpose here, rather than that for which they are used.

It turns out that, in order to meet our purpose of minimising future ranking violations, the order of finish in past races is more important than the distance between runners which drives the RPR awarded in a race.

Runners tend to be more tightly clustered in the finish of Japanese turf races, for instance, than they do on dirt in the U.S. The rankings algorithm makes an adjustment for that.

Before rankings are calculated each week, *TRC Global Rankings* works through the network of rankings provided by RPR and re-calibrates and, if necessary, adjusts the figures in order to create the tightest database of results possible. (Here, 'tightest' is defined as that set of ratings that minimises the variation in each individual horse's set of ratings.) We think RPR do a really good job, but, strictly in a rankings context, a computer can tighten the nuts and bolts.

With this done, the translated or time-decayed RPR, **tRPR**, is formed as a weighted average of each competitor's set of RPR where the weights are determined by

· The recency of the race, because newer evidence is more predictive of near-to-hand success – we use exponential down-weighting of older races

· The number of times each horse contributes to the competitor's total – five individual G1 winners say a lot more than one horse who wins five G1s

Finally, and most importantly, the resulting weighted average is shrunk back towards what we define as 'replacement level' depending on the size of the sample. As explained in detail in section 4.1 onwards, larger samples of success are more reliable than smaller ones; in order to balance the aggregate of achievement with its efficiency, *TRC Global Rankings* uses the simple-but-powerful technique of **Laplace smoothing**. Sections 5.1 – 5.3 contain the details.

3.2 How the rankings use win-loss records

IV stands for 'Impact Value'. This is an oft-used metric in racing analysis and is much better than its more basic cousin winning percentage (often referred to as 'strike-rate'). **IV** improves on winning percentage by taking into account the number of opponents defeated. After all, it is much more impressive to earn a winning percentage of .200 (20% strike-rate or one win in five) in fields of 20 runners than in fields of 10.

To take this into account, IV compares the actual winning percentage achieved by this competitor with the average winning percentage expected given the number of opponents faced. The formula is simple: **IV** is winning percentage divided by the reciprocal of ('one over') average field-size. For instance, if a competitor earns a winning percentage of .178 (17.8% strike-rate) in average fields of 9.3 runners, **IV** is .178 / (1/9.3) = 1.66. Section 5.3 (ii) has another worked example.

The **IV** for each competitor is then given the same treatment as detailed in **tRPR** above:

i) It is weighted for recency and diminished by some function of the number of times an individual horse counts towards a competitor's total;

ii) It is shrunk towards replacement level to equalise for the specific context of each competitor's location and context. See Section 5.1 onwards.

3.3 Fine-tuning the coefficients to meet the objective of the rankings

All the above is technical and gets a further treatment later on, but it is important not to drown in the details. Let's reprise: the mathematics that drives is focused on one objective - to minimise ranking violations (i.e. predict the finishing order of future races by the order of ranking) in all Group or Graded races run before the rankings are reissued in one week's time.

****
****

This is the definitional form of **points**: in any of the four categories of **JOCKEYS**, **OWNERS**, **TRAINERS** or **SIRES**, an entity has a higher (lower) **points**score than another because it is less (more) likely to be the cause of a ranking violation in a near-to-hand race. **Points** are comparable across categories also.

****
****

In this way, *TRC Global Rankings* are formed as the ultimate 'power ranking' – a quantitative assessment of relative competitive strength, scaled so that the numbers have more eloquence than any words. Across time back to 2014, we know they work (or, at least as well as we can possibly make them work) in balancing past accomplishments with future potential, and it is the ambition they are both a classification of the structure of global Thoroughbred racing and a useful tool in projecting the future course of the sport.

3.4 The Modal Country of a competitor

Something may have aroused your interest regarding nationalities in looking through the tables of rankings on the site.

That's right. It said '**Godolphin** (Australia)'. Why Australia, not Dubai?

It's always been one of the most important aspects of the *TRC Global Rankings* that we don't just use data (from the *Racing Post*) but generate it ourselves too. We don't know of anyone else who, at the click of a mouse, can trace the career trajectory of humans involved in the sport, quantifying their influence at any point since 2014.

So, we didn't want to 'waste' one of the fields we display in the rankings table on something you already knew or could easily find out, that **Godolphin** is to the core a Dubaian racing and breeding organisation – “located in the United Arab Emirates” it says on the website.

Godolphin owns horses all over the world, and if anywhere is more closely associated with Europe. But the phrase “it’s been a good week for Godolphin-owned horses” used contemporaneously in different parts of the world could be true or false at the same time. It depends.

An Australian racing analysts might be implying ‘(in Australia)’ when referring to Godolphin-owned horses, whereas a European might variously be implying ‘(in France)’ or ‘(in Britain)’ without explicitly including the national qualifier. So what do we mean by including ‘(Australia)’ in the text above, or listing ‘Australia’ in the column **Modal Country** column in the rankings displayed on the site?

First, *TRC Global Rankings* uses all data, no matter the location, in any row of the table alongside **Name**, which in this case is just ‘Godolphin’. It says ‘Australia’ in **Modal Country** because this is the country in which Godolphin have had the most runners within the ranking period.

You will see that this is data in and, of itself, worth not just detailing but also archiving in the *TRC Global Rankings* annals.

It also side-steps the difficulty of determining national identity in the case of stallions who shuttle or trainers with satellite barns in different countries. Sure, **Mike De Kock** would identify himself as a South African, but it would be misleading to call Hong Kong-based **Richard Gibson** a British trainer or, conversely, **Mirco Demuro** a Japanese jockey.

Instead, we will let the data do the talking, and often it will tell us something interesting.

3.5 Rankings specific to the racing in a certain country

As the above said in 2.1, we will sometimes cast *TRC Global Rankings* using the results of races in only a certain country, a continent even, or just on grass or dirt or for fillies and mares. This isn't going to be something seen every week, but the system is usually sufficiently stable to do this, even with a vast reduction in the size of the sample data.

We actually do this internally all the time to monitor global trends in the sport, which we will bring to your attention when justified.

4.0 The data

TRC Global Rankings use the results of all Group and Graded races ratified by the International Federation of Horseracing Authorities and contained within the database of the Racing Post, which supplies us with the data.

Only races within three years of the rankings date (usually the end of Sunday's racing each week) are considered, with a time-decaying coefficientapplied to weight recent results more heavily than old ones. As with all the moving parts in the rankings algorithm, the coefficient is optimised via millions of back-tests.

4.1 Race qualification; 'Trend' and 'Seasonality' in the rankings

In effect, the weekly **points** scores of a competitor across weeks and months constitutes a time series. Our use of a time-decaying coefficient means that results are lagged across time and diminish in importance as they become further removed from the date on which the rankings are cast.

In a time series, last week's score is obviously correlated with this week's because there is a lot of the same data common to the calculation of both. Because of this, it is necessary to comment on two important aspects of a time series germane to our intention here:

Trend: The total **points** scored by all competitors in all categories is some function of the total races run and the total runners within those ratings. This increased significantly over the last decade, but the upsurge was arrested in 2015 (particularly in South America) and is likely to suffer reduction due to rationalisation in future (**Comment**: *it is somewhat surprising that the situation has been allowed to get so far out of hand.*)

As a result, there is some trend in the level of **points**, but this is minimal and discernible only over multiple years. It would be remiss not to mention this, however.

Seasonality: Some racing countries have a more defined 'close-season' than others, and, even in countries like the U.S., where Graded racing takes place all year, there are periods of more intense competition than others. For this reason, *TRC Global Rankings* exhibit a degree of seasonality.

The inactivity of an entity in the nominal close-season of its **Modal Country** gradually leads to a decrease in **points** because this creates greater uncertainty that past results are reproducible.

Consider the possible reasons for this in the case of a competitor in **TRAINERS**, for example. A new season could mean a new generation of horses, a change in ownership base, a freshly minted stable jockey. *TRC Global Rankings* will, of course, evaluate whether these things result in a change of efficiency as the evidence accumulates, but it pays to temper expectations somewhat in the meantime.

4.2 Backtesting and optimisation

The evolution of the coefficients that control *TRC Global Rankings* started at the beginning of 2014, and all past weekly rankings are archived internally. (We will be drawing upon this data in some accompanying articles each week.) To perform the back-test, a random point is chosen in the past and the coefficients optimised retrospectively to minimise the ranking violations occurring in the results which then followed.

Harnessing the power of modern computing and code (we use R to do the heavy lifting) we can do this millions and millions of times, to the point where we know – to the best of our ability – the coefficients cannot be beaten. It is, however, a dynamic world in which we live and the multitude of factors *TRC Global Rankings* doesn't know or cannot quantify, not to mention the vagaries of human decision-making, make the actual future more challenging to predict than the 'virtual' or 'known' future that backtesting concerns itself with. Again, we are very careful not to overfit the coefficients, as mentioned in 1.0.

5.0 Computation of rankings points and mathematical details

At the first rankings meeting we ever had, a guiding principle was established that hasn't, and never will, change. It serves as the *TRC Global Rankings*mission statement:

****

The rankings **points** awarded should marry the aggregate of success with the efficiency of its achievement. The total number, and quality, of races won should matter greatly, but there should be a debit for profligacy. Lower-profile operators with high efficiency levels should be recognised.

****

There were several mathematical paradigms that could have been used to achieve this trade-off between brute force and efficiency, but the approach we settled on is an adaptation of Additive Smoothing, itself an extension of a famous idea in probability, Laplace's Rule of succession.

5.1 Introducing Laplace smoothing

Say you know nothing about a team in a sport where games are settled with no ties possible. What is your estimate of their winning percentage? It has to be 50%, of course, or .500 as it is usually expressed about teams in American sports leagues.

With this as your prior expectation, now you are told the team lost its first game of the season. With a record of 0-1 now realised, the league standings record its winning percentage as 0% or .000. But your estimate of its future success-rate should clearly be higher than this. Nearly all sports teams lose a game at some point in the season, after all.

With virtually no information about the team to guide you, making the best guess is entirely theoretical. Motivated either by estimating the probability of the sun rising, or finding the likely location marker on a craps table, Laplace suggested that, abstracted from other information, the best guess about the team's winning percentage for the rest of the season is achieved by simply adding one to the team's win total and one to its losses. So, after a sole loss, the best guess for our team's winning percentage – calculated as wins / (wins + losses) – is 1+0 / (1+0 + 1+1) = 1/ (1+2) = 1/3 = .333

Probability theory has moved on a little since the beginning of the 19th century, and we now understand that what is going on here is better known as "the estimation of a proportion by a non-informative prior", which is part of the increasingly popular discipline of Bayesian inference, first developed well before Laplace by the Reverend Thomas Bayes.

All that the Rule of Succession does is start by assuming the team above is already 1-1 before a ball is kicked or thrown or whatever. But, why choose one win and one loss and not, say, six wins and six losses?

The answer is that, abstracted from any other information, we didn't have any reason to choose another number but one. Adding one win and one loss to the denominator of the calculation of win percentage – wins / (wins + losses) - makes sense because it avoids 'division-by-zero error' while avoiding biasing the estimate towards either more wins or more losses.

If, instead, we add six wins and six losses to the record of the team above which loses its first game, our estimate becomes (6+0) / (6+0 + 6+1) = 6/13 = .462.

This new estimate of .462 is much greater than the .333 we calculated when starting with only one win and one loss. What's going on here is regression towards the mean of .500. The more wins and losses we see before the seasons starts as our prior estimate of success, the less difference one extra 'real' game makes when we add on a single loss.

In a real sports league, like the NFL in the US, we know the long-term variance in the end-of-season winning percentages of teams, which is a function of the ratio between skill and luck in determining results, as this explanation shows. So, rather than choosing a prior of one win and one loss before the season starts, the best estimate turns out to be adding about five wins and five losses. Before we see how this applies to *TRC Global Rankings*, let's work through an estimate.

After four games of the 2016 NFL season, the **Atlanta Falcons** had won three games and lost one. This is a winning percentage of 3 / (3+1) = .750. You might hear some analysts refer to the Falcons as "being on pace to win 12 games" for, if they maintained their winning percentage over the NFL's 16-game regular season, they might be expected to win 12 and lose 4.

This expectation of .750 for the remainder of the Falcons season is generally too high, however, and a better estimate can be achieved by adding five wins and five losses to their record of 3-1, making the best estimate the result of the calculation (5+3) / (5+3 + 5+1) = 8 / (8+6) = 8/14 = .571

You should be able to see that, using the pre-season prior of five wins and five losses, the Falcons' current winning percentage of .750 is regressed towards the mean of .500 for all teams in the NFL and becomes .571. With 12 games left to play, the best guess is the Falcons win .571 x 12 = 7 which, added to their existing wins and losses in the bank, would see them finish with 10 wins and 6 losses – not 12 and 4.

(A salutary lesson of regression towards the mean is provided by the Falcons 2015 season, which was 4-0 over the first four games and 4-8 over the final 12 for a total of 8-8 or .500.)

****
****

So, additive smoothing (named after Laplace) adds n wins and p losses to a team's record to smooth (or regress) a team's actual win-loss record towards the league average, so that the best estimate for the expected winning percentage going forward is formed by n+W wins and p+L losses, and is (n+W) / (n+W + p+L)

****
****

5.2 Extending Laplace smoothing using the concept of replacement level

In the example above, we conceive of Laplace smoothing as placing a non-informative prior over a team's record, expecting the team to have a .500 record in the light of any other expectation.

Another way to look at the addition of five wins and five losses to a team's actual record is to think of it as rewarding teams with a greater sample size of winning performance than those who have won only a few games. Consider the effect of a Laplace smoother of n=5, p=5 (five wins and five losses) to two different unbeaten teams, the first after three games and the second after eight games.

For the 3-0 team, the expected ***or smoothed*** winning percentage is (5+3) / (5+3 + 5+0) = 8/13 = .615

For the 8-0 team, the expected***or smoothed*** winning percentage is (5+8) / (5+8 + 5+0) = 13/18 = .722

We would rank the second team much higher than the first simply because there is more evidence to go on. Both teams are unbeaten and therefore have a winning percentage of 1.000 (100% record), but the second is highly likely to be the better team.

Can you see where this might be leading for ranking jockeys, owners, trainers and sires?

For two jockeys or owners or trainers or sires who have the same, above-average **win percentage**, the higher-ranked entity should be the one with the greater **sample size**. After all, a lot of trainers can send out five runners in Group races and have one winner, but only a top-ranked trainer can repeat this one-in-five strike-rate in a hundred races. In the end, regression towards the mean (or "statistical gravity" as it could be called) brings down all but the truly elite.

Now, in the NFL example, the method of adding the same number of wins and losses to a team's record was called a "non-informative" prior because it treated all teams the same and weighted them to the league-average rate of wins and losses. In our setting, we know an absolute ton more information about the competitors in each of our four categories of JOCKEYS, **OWNERS**, **TRAINERS** and **SIRES**.

Most importantly, unlike the non-informative prior, we know that competitors with lots of runners *tend* to do better in future than those with few. The hierarchy of competitors, from the powerhouse to the one-horse show, *tend* to have got where they have by being good at what they do. (It is massively important here to remember what *tend* really means in statistical language – there is an 'error term' or residual in the equation – and that it means that exceptions to the rule being postulated are not just inconvenient truths, but are all part of normal expectation: there are brilliant operators among those barely respected as belonging to the elite of the game, while others don't quite punch their weight.)

Using Laplace smoothing, *TRC Global Rankings* in effect place the aggregate of success in the context of the efficiency of its accomplishment. More runners mean more thrust against statistical gravity, but only the truly efficient reach the stars.

So, the *TRC Global Rankings* algorithm chooses a prior designed to equalise all the different circumstances in which jockeys, owners, trainers and sires find themselves. Regression takes place not towards the mean, but towards what the rankings algorithm calculates is a suitable replacement level for each category.

Replacement level can be thought of as that rate of achievement that is best described as 'meh'. It is usually a long way below the mean; it is that rate of success which a replacement competitor could easily achieve. In TRC Global Rankings terms, we calculate replacement level carefully for each category.

5.3 A worked example of ratings calculation, showing how Laplace smoothing works

Let's go back to the first-ever week of rankings in January, 2014, and work through the calculation of the rating for **War Front** in the **SIRES** category. Here are the steps in the process:

i) All *Racing Post Ratings* are revised and rescaled for rankings purposes. We use adjustments according to the surface and the country to make the order of finish more important than the distances between runners, then connect the thousands of performances together as tightly as possible using mathematical methods;

ii) We transform the resulting RPRs for War Front to reflect their recency and to diminish multiple RPRs earned by the same horse, so that five individual winning horses count for (a lot, as it happens) more than one horse earning winning five times;

iii) The resulting **tRPR** for War Front's sample of 137 runners was initially 97.7 and their Impact Value (**IV**) was 2.23. The latter figure indicates that War Front's stock won 2.22 times more often than random, considering the size of the fields in which his runners competed. (For the record, his strike-rate was 23.4% and the average field-size 9.52, so the calculation for **IV** is .234 / (1/9.52 = 2.23.)

iv) The figures 137 runners, **tRPR** 97.7, **IV** 2.23 - are stored. (1)

v) Next, the algorithm determines 'replacement level' for a competitor in the **SIRES** category in that week. It had already iterated over millions of pre-calculated **SIRES** rankings to find the optimal replacement level by trial-and-error.

At that point, replacement level for **SIRES** was 142 races of RPR 79.3 and 134 races of **IV** 0.79. Let's go over what that means, first for **tRPR**.

Before any of the data for **SIRES** is taken into account, the system calculated that a replacement-level stallion in the **SIRES** category would have had 142 races and a **tRPR** of 79.3. Being one of the world's best stallions, War Front is miles better than this, of course. His stock had 137 runs and a **tRPR** of 97.7.

Now, the algorithm knows from testing all sires with sample-sizes like War Front that their statistics tend not to be 100% reliable in forecasting the next week's results. There is just too much randomness (variance) in racing. This is where Laplace comes in, and his method of additive smoothing is simply a function of combining the replacement sire's data with War Front's.

So, War Front's smoothed **tRPR** is his actual data 137 runs * 97.7 plus the data of the replacement sire 142 runs * 79.3 all divided by the total number of runs of War Front and the replacement sire which is 137 + 142.

Again,

[(137 * 97.7) + (142 * 79.3)] / (137 + 142)

which equals 88.3

Therefore, War Front's actual **tRPR** of 97.7 is shrunk to a smoothed **tRPR** of 88.3 (2)

Can you see the idea here?

Say, War Front's sample-size was actually twice the real one, 274 runs at an average **tRPR** of 97.7, instead of 137 runs. To have sustained this level of elite performance for twice the number of runners is a much, much more impressive performance.

In this case, the calculation for the smoothed **tRPR** would be:

[(274 * 97.7) + (142 * 79.3)] / (274 + 142)

which equals 91.4

Now, War Front's actual **tRPR** of 97.7 is shrunk to a smoothed **tRPR** of 91.4 which is much higher than (2) above.

The larger the sample-size, the less smoothing takes place.

That is why you will see that the rankings tend (that word again) to be headed by the **Coolmore Partners** and **Chris Waller**s of this world. But the smaller men, women and sires of this world can still penetrate the upper echelons, despite Laplace and his infernal smoothing, if their average is high enough.

Using Laplace smoothing, we can equalise the variance of every different sample size in the **SIRES** category, using the original **tRPR**. The resulting smoothed **tRPR** (which appears in our ranking tables) is then directly comparable for every competitor in every country and every sample-size.

5.4 Smoothing Impact Value

The final, smoothed **tRPR** is one part of the calculation for ranking **points**. It's exactly the same deal with the other component, **IV**.

Let's continue the calculation for War Front:

War Front: 137 runners **IV** 2.23

Replacement sire: 134 runners **IV** 0.79

The calculation is (137 * 2.23 + 134 * 0.79) / (137 + 134)

which equals 1.52

Therefore, War Front's actual **IV** of 2.23 is smoothed to an **IV** of 1.52 which is comparable with every other competitor of whatever sample-size in **SIRES**.

5.5 The final calculation of ranking points

After smoothing, War Front's **tRPR** is 88.3 and his **IV** is 1.52. These are combined in the optimal way to form **points**.

The precise methodology here is proprietary, and is different for each category, but you should be able to see how **tRPR** and **IV** combine by running your eye down the table or using numerical methods.

And that's it. **Points** determine **RANK** and the system waits for the races to be run to test its predictions …

James Willoughby, October 2016