European Journal of Operational Research 247 (2015) 641–647

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Innovative Applications of O.R.

What is a good result in the first leg of a two-legged football match?

Ramón Flores [a], David Forrest [b,*], Cesar de Pablo [c], J. D. Tena [d,e]

[a] Universidad Carlos III, Deapartamento de Estadística, Avda. Universidad Carlos III 22, 28270 Colmenarejo, Madrid, Spain
[b] University of Liverpool Management School, Chatham Street, Liverpool L69 7ZH, UK
[c] Daedalus, S.A.. Edificio Vallausa II, Avda. de la Albufera 321, 1ª Planta - Oficina 10 28031 Madrid, Spain
[d] Universidad Carlos III. Departamento de Estadística. C/Madrid 126, 28903 Getafe (Madrid), Spain
[e] Università di Sassari, Dipartimento di Scienze Economiche e aziendali, Via Torre Tonda 34-07100 Italy

ARTICLE INFO

Article history:
Received 11 August 2014
Accepted 28 May 2015
Available online 6 June 2015

Keywords:
OR in sports
Football
Two-legged matches
Home advantage

ABSTRACT

The most important pan-European football tournaments include ties where two clubs play each other over two matches and the aggregate score determines which is admitted to the next stage of the competition. A number of stakeholders may be interested in assessing the chances of progression for either of the clubs once the score of the first match (leg) is known. The paper asks what would be a "good" result for a team in the first leg. Employing data from 6,975 contests, modelling reveals that what constitutes a good result has changed substantially over time. Generally, clubs which play at home in the first leg have become more likely to convert any given first-leg result to eventual success. Taking this trend into account, and controlling for team and country strength, a probit model is presented for use in generating probability estimates for which team will progress conditional on the first-leg scoreline. Illustrative results relate to ties where two average teams play each other and to ties where a relatively weak club plays home-first against a relatively strong club. Given that away goals serve as a tie-breaker should aggregate scores be equal after the two matches, the results also quantify how great the damage is when a home-first club concedes an away goal.

1. Introduction

The two pan-European competitions organised by UEFA (the Champions League and the Europa League) are the most lucrative club football tournaments in the World. For the organisers, they generated commercial income, principally from the sale of television rights, of more than €1.6b in 2013–14 (www.uefa.com). Their structure has varied over time, no doubt in order to increase this commercial income further, but currently features both 'group' and 'knock-out' stages.

The group stage comprises mini leagues in which each club plays each other club at home and away to determine which two in the group will proceed to the next phase of the competition. The knock-out stages, including the rounds leading to the Final, are organised on a straight elimination basis such that pairs of clubs play each other twice (once at each of the home stadia) and the aggregate score over the two matches ('legs') determines which will survive to the next round. If the aggregate scores are equal, the first tie-breaker is the number of away goals scored by each club. If this still does not settle the issue, resort is made in the second-leg, first to 30 minutes extra time and then, if necessary, to a penalty shoot out.

Operational researchers have investigated a range of issues in sport (Wright, 2009, 2014) including proposing models for the evaluation of performance after a contest (Fried, Lambrinos & Tyner, 2004) and for forecasting final outcome as an event unfolds (Klaasen & Magnus, 2003). This paper touches on both these themes. It seeks means of offering guidance to interested parties once the first-leg score in a European tie has been determined, in terms of how satisfactory the result was for a club and what the prospects are of the club advancing to the next round.

A probabilistic assessment of which club will progress is likely to be useful in decision taking by a variety of stakeholders. For example, the organisers of the competition will likely wish to assign the best qualified referees to second-leg ties where the outcome of the tie is still most in the balance. Broadcasters in many markets have to choose which of several matches on the same night to show on television or assign to their principal channel; the degree of suspense remaining over which team will survive to the next round will be relevant. Supporters of the club which played at home in the first-leg must decide whether it is worthwhile to make an expensive international journey to attend the second-leg; they may wish

* Corresponding author. Tel. +44 151 795 0679.
E-mail addresses: rflores@est-econ.uc3m.es (R. Flores), david.forrest@liv.ac.uk (D. Forrest), cdepablo@daedalus.es (C.d. Pablo), jtena@est-econ.uc3m.es (J.D. Tena).

http://dx.doi.org/10.1016/j.ejor.2015.05.076
0377-2217/

to conserve their funds if they would be travelling for an almost lost cause. Coaches and players have to decide how much effort to place into the second-leg; here again it matters whether the second-leg is almost a 'dead rubber' or whether everything remains to play for (for example, a coach may rest key players to conserve their energy for upcoming domestic matches where the probability of advancement in Europe is either very high or very low). Finally, bettors and bookmakers may wish to engage in wagering on the ultimate outcome of the tie and therefore need to assess the probabilities at the mid-way point.

To all these actors, the model we propose has potential value in decision making in the time between the first- and second-legs. For coaches, the model may also have utility prior to and during the first-leg. In formulating strategy, a coach will need to think about what would be an appropriate target result. Possible strategies, ranging from very defensive to very offense-orientated, carry different degrees of risk. A coach may, for example, eschew further risk (and switch to emphasising defence) once his team has established a goal superiority which, if maintained until the end of the first match, would give a satisfactory probability of progression to the next stage of the competition. In offering a means to estimate probabilities of progression, the paper therefore serves also as a contribution to sport analytics, the application of statistics and operational research to strategy formation in sport.

We employ a data set we assembled which includes all 6975 two-legged ties played in the Champions League and the Europa League, and their predecessor competitions, from the introduction of the away goals rule (for settling ties where aggregate scores are level) in the 1960s to the end of season 2012–13. Controlling for team strength, we investigate the significance of different first-leg scores for the chances of either team progressing to the next phase. What constitutes a favourable outcome for a club will prove to have changed over time and we will therefore present illustrative probabilities based on estimation of a model using only the final sixteen years of the data period.

2. Data

We first collected data on all past results in the histories (1955–2013) of the UEFA Champions League and the Europa League and their predecessor competitions.

The source data archive was that provided by the Rec. Sport Soccer Statistics Foundation (www.rsssf.com), which has been compiled and curated by volunteers over a long period.

This data archive offers an invaluable resource; but extracting information to be used in data analysis and for generation of predictor variables involving clubs' past performances is not straightforward and presents many challenges. For example, the structure in which the data are held varies over matches and club and country names are often entered inconsistently (one club name, Ferencvárosi T.C. appeared in eight different variants). Our process for building our own data base from the source data base included (i) crawling the files at www.rsssf.com, (ii) using Information Extraction techniques to obtain information on match and tie outcomes, (iii) using Information Integration to address the problem of inconsistent club names, and (iv) creating covariates based on performances in previous seasons of European competition prior to the subject tie taking place.

From the data set we constructed, we use in modelling only ties played under the away goals rule for determining progression to the next round when the aggregate scores are level (in early years, either a coin was tossed or a third game was organised in a neutral country). This means that we do not include in subsequent analysis any ties played before 1965 and only a proportion of those which took place between 1965 and 1970. This still permits consideration of 6975 two-legged ties.

3. Measuring team strength

Our focus was on probabilistic prediction of the outcome of a tie conditional on the score in the first-leg. But, clearly, the probabilities for a given score will vary with the strength of the two clubs.

Page and Page (2007) presented an unconditional model for forecasting the outcome of a tie (i.e. to be used prior to the first-leg score becoming known), employing a large data set which had several years in common with ours but which combined contests with and without the away goals rule in force. In constructing their control variables for team strength, they did not seek to capture domestic records of participating clubs (by metrics such as win-ratios or league position) because there is so much heterogeneity across Europe in the quality of opposing teams faced by clubs in their respective domestic competitions. Instead, they based their 'ability' variable directly on the official UEFA club ratings used in seeding within its competitions.

Currently the UEFA rating is an unweighted average of the number of 'points' won by a club in European competitions in each of the preceding five seasons, where points are earned according to its match results and the phase of the competition to which the club survived. For example, taking part in the group stages of the Champions League gives a club 4 points; it earns 2 points for each win and 1 for each draw; and it receives 5 more points for making it from the group stage to the last 16. In addition, a club's rating also includes 20% of the value of a similar index calculated according to the performances of all the clubs from its national league in the preceding five seasons.

From the perspective of using it in a statistical model, the official index has obvious weaknesses. It treats performances over the previous five seasons equally whereas information from more recent seasons may be more relevant to predicting the outcome of a tie. Further, the relative weights given to performances by the club itself and by all clubs from its national league appear to be arbitrary. Therefore we created our own indices to capture strength though based on the same informational input as the official ratings.

We use separate ratings of club and country strength as covariates in the statistical model to avoid the arbitrary weightings used in UEFA's calculations. We also give more weight to more recent years (the most recent year has the weight one, the preceding year a weight of one-half, the year before that of one-third, and so on). Of course, it could be argued that imposing a particular rate of decay of weights is itself arbitrary. Therefore we re-estimated the models below, replacing the ratings of club and country strength for the two teams using separate indices for each of the five years on which the strength variables are based. This has the virtue of allowing the data to determine weights and the number of observations is comfortably large enough to accommodate the increase in the number of parameters to be estimated. However, the gain achieved, as measured by forecasting efficacy, was always marginal (for example varying the Brier score by no more than 0.1). Therefore, in the interests of parsimony, we report only results from a model where clubs' strengths are measured by single indices of each of club strength and country strength.

Points for inclusion in the strength indices were calculated as follows.

Whenever a club was observed to have reached a round with n clubs remaining in the competition, it was awarded $1/n$ points for the particular season. This award was doubled if the competition was the Champions League (or its predecessor, the European Cup) as opposed to the lower status Europa League (and its predecessor tournaments).

Because there were changes in the number of rounds or phases over time, a club's score (total points) for any particular year was then normalised by the total score calculated across all clubs in the competitions that season. Finally, this normalised score was multiplied by 1000 (to avoid working with very small numbers). This was our club score for a single season. The club's strength at season t was the weighted average of its scores between seasons t-5 and t-1, as described earlier.