Competitiveness in Soccer Leagues: An ordinal time series approach with evidence from the Premier League 1993 to 2014

Jones, Calvin and Jordan, Declan

Cardiff Business School, University College Cork

December 2014

Competitiveness in Soccer Leagues: An ordinal time series approach with evidence from the Premier League 1993 to 2014

Calvin Jones[1] & Declan Jordan[2]

[1]Cardiff Business School

Colum Drive

Cardiff CF10 3EU

Jonesc24@cf.ac.uk

[2]School of Economics, University College Cork

College Road

Cork

d.jordan@ucc.ie

Abstract

This paper presents a novel approach to measuring competitiveness in sports leagues and tests the measure using evidence from the Premier League from 1993 to 2014. Traditional approaches to measuring the level of competitiveness have focused on win percentages and/or the dominance of subsets of teams over a number of seasons. The paper argues that these measures are not fully reflective of the core outcomes that teams and fans use to measure performance. An ordinal times series approach is proposed as an alternative measure of competitiveness. Applying the approach using data on end-of-season positions from 1993 to 2014, the paper finds a marginal decrease in league competitiveness over the period. The flexibility of an ordinal analysis approach is further demonstrated by considering the relative performances of subsets of teams based on location. The paper finds that North Eastern teams have suffered a long-term decline in performance on average while London teams are the best performing, and that their relative performance has also improved over the time period.

Keywords: Competitiveness; Competitive balance; Ordinal Time Series; Sports economics

Competitiveness in Soccer Leagues: An ordinal time series approach with evidence from the Premier League 1993 to 2013

Introduction

Just as traditional economics extols the virtues of competition in market economies, it is now widely accepted within the sports economics literature that competitive balance is a very desirable feature of sports leagues, from the perspectives of fans, clubs, and governing bodies. The concept of competitive balance has been at the centre of sports economics literature almost since the first published works in the area (Rottenberg (1956); Neale (1964)). Zimbalist (2002, p. 112) notes that the importance of competitive balance in sports leagues derives from fans' preference for uncertainty of outcome. Clubs also favour competitiveness in a sports league as the demand for their tickets and merchandise depends on fans anticipating competitive games and leagues. In reality clubs, and their fans, are likely to favour very slight competitive imbalance that favours their own team.

Sports economics is not only focussed on sporting outcomes. A reading of the sports economics literature reveals the potential for this microcosm of society to inform, and be informed by, the way we understand wider economic and social relationships. For example, there is a sizeable literature on the use of sport to inform business decision-making and strategies (for a small sample see Keidel (1985); Gilmore and Gilson (2007); Bridgewater, Kahn, and Goodall (2011); Haskins (2013)) There is also a strand of literature treating teams as economic-competitive agents – effectively firms – and assessing the ability of such a conceit to uncover truths about the ways in which 'competition' works under different condition sets (determined by league rules or economic regulation), with different resource allocations, and differences in concentration levels and dominance of subsets of firms or teams (see, for example, Pinch and Henry (1999); Owen, Ryan, and Weatherston (2007)). Indeed, Palacios-Huerta (2014, p. 3) argues that soccer can be used to help economists understand economic processes because "sports are in many ways the perfect laboratory for testing economic theories".

It is interesting that given the concept of 'league as market' and 'teams as firms' much competitiveness literature downplays the longer term outcomes and characteristics of teams themselves and leads firstly to the use of somewhat flawed (or at least incomplete) measures of success and competitiveness, and secondly to the narrowing of the scope of inquiry and the closing of a number of interesting avenues. This paper argues that an alternative measure of competitiveness, an Ordinal Time Series (OTS) approach using end-of-season league position, provides a measure that is more closely aligned with the season objectives of teams. The next section discusses the widespread use of win percentages as a basis for measuring competitiveness within seasons in order to contrast this with our OTS approach (described in Section 3). To demonstrate the benefit of the OTS approach, it is then used to estimate the trend in competitiveness in the English Premier League (EPL) since its inception in 1992 and, then, the relative performance of clubs based on their geographic location over the same time period.

Measuring competitiveness in soccer leagues

Most measures of league competitiveness use win ratios, points accumulation or points spreads to measure competitiveness ((Fort and Maxcy (2003) and Bowman, Lambrinos, and Ashman (2012)). Szymanski (2001, p. 76) and Owen and King (2014) acknowledge that the standard deviation of win percentages has become the most widely used indicator of competitive balance in the US literature and is also used most often to measure competitive balance in European sports leagues, including soccer.

The prevalence of win percentages as a measure of competitiveness may reflect the importance of the US in the sports economics profession and the corresponding focus on US sports. Generally, US sports do not have drawn games and organise themselves in conferences and play-off structures which makes win percentages a direct measure of team performance within a season. Dobson and Goddard

(2011, p. 127) note that the theoretical frameworks underpinning the drivers of competitive balance in sports leagues have been developed mainly by US economists. This has meant that the features of US team sports leagues have dominated development of the theoretical literature in this area. It is also likely to have influenced the development of methodological and empirical approaches.

This measure, however, is weakened by the fundamental point that winning matches and, as a result, accumulating points are largely a *means to an end* (with the exception perhaps of important derby matches). This is *unlike* the firm, where the extent of dominance or market share is of fundamental importance to the central outcome (profits). Winning matches is, of course important from an emotional perspective, but most soccer teams, owners and fans would favour higher league positions over win percentages or number of points. For example, English soccer fans would favour a fourth place (hence Champions League) finish with a 60% win rate compared to a fifth place finish and a 65% win rate. This is not an abstract situation. In the 2003/04 season Newcastle United finished in fifth position and qualified for the UEFA Cup with a win percentage of 34.2%. The sixth placed team, Aston Villa, won 39.5% of their games, but did not play in European competition the next season. Also in the 2013/14 season Norwich City would gladly have traded their win percentage of 21% with West Bromwich Albion's 18.5% if it meant they too could have avoided relegation.

Szymanski (2001, p. 76) notes that, despite the number of drawn games in soccer, win percentages are a reliable indicator of success since it is "closely correlated with the more usual measures of success such as league position (correlation coefficient 0.91)". However, win ratios and points accumulated remain indirect indicators of team performance relative to that team's core objective, and are used even where the direct variable – the ordinal, ranked position of the team in the league at the end of season – is readily available over a long period. Since league position is available for most soccer leagues (where there is an absence of conference structures and post-season playoffs, such as those present in US professional sports, which would distort comparison) it is possible to use this as an alternative measure of competitiveness.

There is also a limitation in win percentages as a proxy for competitiveness across a full league campaign. Wins (and points) cease to matter as much when the ordinal outcome is unaltered or variation is less important. This is the situation when the league is won, or play-offs, European qualification or relegation has been decided. In this case team management may start to plan for the next season, team ownership may start to plan for new management or players simply expend less effort.

A focus on finishing position enables us to re-cast our approach to understanding competitiveness. We now have a single, bare metric with a fixed relationship set to other teams in the league, having discarded proxy or substitute data. Here, then we can focus more on the variability of finishing positions over *time* as a measure of competitiveness.

Our use of this measure has implications for how competitiveness is considered. Consider a situation where the end-of-season finishing positions remain unchanged from one season to the next, but the league winners significantly increase their points difference to the second place team. While a win percentage approach would find that there has been a reduction in competitiveness between the two seasons, an approach using end-of-season finishing positions would indicate no difference in competitiveness between the seasons. Although the former approach is one appropriate measure of competitiveness, the latter may be more insightful if the relative competitive position within the league is expected to affect long-term club performance more than the points difference between the teams in these positions. For sports teams, we argue finishing position may matter more than other metrics for long-term success as well as short-term fan satisfaction. Points spreads and win ratios do not matter if their variations do not affect the 'success' outcomes of the team *in the long run*. For example, FC Barcelona's championship winning fifteen point margin over Real Madrid in the 2012-3 season did not significantly increase their likelihood of winning in 2013-4 because that points difference led to no more resources available for the following season over and above those associated with winning the Championship by a single point.

There is support within the literature for the view that league position is an important metric of club success. In his seminal paper, Rottenberg (1956, p. 246) asserts that a club's attendance is a function of its local market size, the size and access of its stadium and its "average rank standing…during the season in the competition of its league". Whitney (1988) considers whether win percentages or championship prospects influence fan interest and finds that demand for tickets in US baseball depends in part on championship winning prospects.

A number of papers focus on finishing position as a metric – although this is typically restricted to championship wins or high finishes. For example, Goossens (2006) looks at top 3 finishes in European soccer and qualification for the lucrative Champions' League. However, there is value in developing a whole-league measure. A cursory examination of the attendance figures in any sports league demonstrates of course that fans continue to demand tickets for those clubs that are not likely to win the championship or achieve a high finish. The utility of fans and clubs may not be dependent on the team's championship winning prospects but rather on expectations of finishing position in a given season. Michie and Oughton (2005) report the results of a survey of English soccer fans that states that 89% believe that only five clubs have a "realistic chance of winning the Premier League" in a given season, but still only a minority agree that the league is "boring and predictable" and just over one third are dissatisfied with the level of competitive balance. This indicates that measuring competitive balance by focusing on the finishing positions of only those at the top of the league may not reflect the objectives and expectations of clubs and fans.

In summary then it is notable that league position has been largely downplayed in the literature on competitive balance in favour of measures based on win percentages or the dominance of high finishing positions by particular teams. However an examination of where teams end up at end of a season may be a useful complement to measures based on where the outcomes of matches.

An Ordinal Time Series measure of league competitiveness

This paper argues that the year-on-year change in teams' end-of-season league position has a strong conceptual justification as a measure of league competitiveness. Our measure has an additional benefit in that the actual league finishing position (effectively the first order of our time-derivative variable) is specific to individual teams and hence can be directly related to team attributes that might be the underlying drivers of competitive balance (an exploratory example of such an analysis is provided in Section 5 below).

We might surmise a dataset of n teams over t years, with end-of-season position, R_t, our first order team performance metric. Where there is perfect competitive balance in a league, that is where there is an unpredictable outcome, it would be the case that end-of-season position would be independent of the end-of-season position in the previous year. That is, R_t would not be associated with R_{t-1} and so would effectively be random. This is unlikely to be the case, since the rewards associated with end-of-season position would be expected to influence a team's ability to attract better playing talent in subsequent seasons (25% of a team's revenue from the Premier League is based on its end of season finishing position).

One appropriate measure of a league's level of competition is the volatility in end-of-season position for teams from season to season. This facilitates ranking and comparison of different leagues in a given season and comparisons over time (allowing of course for different league sizes and the impact of relegation 'break points'). Developing a 'holistic' hierarchical measure is, however, not straightforward. Prior uses of finishing position have focussed on specific *positions* – how many different winners of the league or the number of teams in the top ranked places over a given period for example – and while this is limited in scope, it does avoid the knotty problem associated with relegation and promotion. Where teams are relegated out of (or promoted in to) a league each season, R_{t-1} is not available (in that league) for some teams. There are a number of ways to resolve this issue.

1. Restrict the analysis to the set of teams that is unchanging within the time period of analysis.
2. Restrict the analysis to leagues that do not allow promotion and relegation.
3. Extend the analysis to allocate ordinal positions to lower leagues.

There are limitations in each of these approaches. The difficulty posed by the first solution is that it results in selection bias and significantly reduces the population size. For instance, only 7 teams have appeared in the Premier League in every season since it began in 1993.

The second solution would restrict the analysis to very few soccer leagues, as the pyramid structure is integral to most soccer associations. There are also very few leagues in other sports that do not have promotion and relegation, with the exception of US sports leagues. Szymanski and Zimbalist (2006) discuss the different traditions in US and European sports and note that US sports tend to operate on a franchise basis without promotion or relegation. However the effect of pro-competitive interventions such as salary caps and draft picks mean an analysis using these leagues is not 'policy off': there are explicit pro-competitive interventions by regulators that do not occur in soccer leagues outside the US. Additional issues would arise from the conference structures of US leagues and from play-off periods, where we would be faced with multiple and non-ordinal finishing positions (e.g. semi-finalists).

The third solution presents a higher level of complexity in data collection and in conception and analysis. This is because, effectively, the results would relate not only to the top-tier league but also to the system of leagues in a given country, or (as what follows will show), the subset of teams in the national leagues that have shown themselves to be capable of playing at the highest level.

In this paper we use the third approach, examining a league that allows relegation and promotion and focussing on the English Premier League, including the 44 teams who have appeared in that league for at least one season between 1992-1993 and 2013-2014. Teams that were relegated from the Premiership were awarded a position based on their finish in the overall Football League, ie. R + 20 (or + 22 in the first two seasons of the Premier League). Our analysis is facilitated by a national league structure in England that is fully hierarchical so that each team has a unique ordinal position within the national leagues each season.

Our ordinal time series (OTS) analysis has a limited history in the social sciences, with some application in management (see Timothy W Ruefli (1990)), but most discussions are theoretical or in computing and hard/medical sciences (Bandt (2005) and Bartimote-Aufflick and Thomson (2011)). Here we follow the work of T. W. Ruefli and Wilson (1987) who used OTS to examine competitiveness and industry structure in US transportation 1954-1982, largely as a demonstration of OTS benefits. For these authors, OTS was useful as it firstly, obviated the need for consideration of a number of factors that were inherent with cardinal economic analysis over time and between firms, including relative price baskets for different firms and reflation across time for key metrics. They also argued ordinal analysis did away with problems due to dirty or incomplete data, and inherently displayed information about other firms in the rank of the target firm that cardinal data did not – i.e. if PanAm slip 2 places in the ranking, 2 other firms have improved whereas an improved profits performance by PanAm may still be worse than all others in the industry.

Many of these benefits are not relevant to sports leagues, where data are not financial, and of very high quality, but here we have a strong conceptual reason to include ordinal measures (as outlined in earlier Sections) and the metrics that T. W. Ruefli and Wilson (1987) used are applicable here.

OTS measures can be used to understand competitiveness in a number of ways. Movement in ranking over time might help identify;

- Changes in competitiveness in one league over time,
- Relative performance of different sub-sets of teams (as the average rank of that set moves up or down over time),
- Relative volatility (competitiveness) between different leagues,
- Persistence in performance in proximate years.

This paper does not explore all of these uses but, to illustrate the method, presents an analysis addressing the first and second points.

Applying OTS: Volatility and Competitiveness over Time

The competitiveness C of the set of teams N across years t = 1 – t is given as

$$C(N_t) = \Sigma_1^n[|R_t - R_{t-1}|]$$

Thus, year-pairs which display greater movement in rankings between them are considered more competitive.

[Figure 1 Here]

Figure 1 suggests that there is some, evidence of the Premiership becoming a *little* less competitive between 1992-3 and 2013-4, albeit marginally – teams are moving slightly fewer places between proximate years later in time. The average team movement is at a maximum of 7.3 places in 1995-6, and there is a period of lower volatility from 2006-7 to 2011-2, including the low of 4.6 average places moved (the series average is 5.8).

The finding of a marginal decrease in competitiveness is consistent with Goossens (2006) who, using a measure of trend dominance which is the number of teams finishing in the top three league positions over a five season period, finds that there has been a decrease in competitive balance in the English league since the mid-1990s.

Similarly, Curran, Jennings, and Sedgwick (2009), by constructing a Top 4 Index, which measures the proportion of top 4 positions occupied by the four most successful teams over a given period, find that competitive balance in the top division of English soccer has decreased relative to the levels observed in previous decades.

As noted earlier, our results include data from outside the Premiership – including teams that are relegated or promoted to/from the Football league during the period, and with these data points and teams likely to display higher volatility than Premiership-only changes in position (the range of potential finishing positions is greater). By restricting the pair-analysis to pairs where both positions are less than 21 we restrict our analysis to competition within the Premiership itself – excluding all instances where a target team were promoted or relegated within that pair of years. Our 'N' of included pairs reduces from 924 to 353 within-Premiership rank changes.

Here we see a similar pattern to that for all 'Premiership-appearing' teams, albeit now of course with far lower range of potential changes in position year to year (theoretically greatest at 19 places, from champion to last). Whilst the pattern is similar to all teams, within Premiership competition seems to have declined at a faster rate – a monotonic trend of -0.063 compared to -0.047. Whilst this fits with the wider meme – that of the 'locking up' of the top four Champion's League places by a select sub-set of clubs – it should be noted that competiveness overall increased in the first few years after the shift to a four-club qualification system (2001-2) hinting, perhaps, at lagged effects in consequent financial concentration and club power.

[Figure 2 here]

Whilst there is an overall declining trend, the measure is quite volatile. For example, the period 1998-2001 was characterised by low competitiveness, but higher competitiveness in the period 2002-2006, post-Champions League qualification changes.

Applying OTS: Assessing the Performance of Different Team Types

Another use of ordinal analysis suggested by T. W. Ruefli and Wilson (1987) is analysing sub-sets of firms (in their case airlines, trucking and rail subsectors of US transport) to understand changes in

relative performance over time. A similar analysis may also be undertaken for subsets of teams in sports leagues. Teams could be classified by a variety of characteristics, for example ownership model, fan base, financial structure, to examine the success, in a given season or over time, of different team sub-sets. The performance P of a sub-set S in season t, measured by the teams' average end-of-season league position, would be given by:

$$P(S_t) = \Sigma_1^s R_t / N$$

R_t is the ordinal rank and N the total number of teams.

In many cases of course a team's classification might change within the time period – for example, a changed ownership structure or financial status/strength – and the classification set would then be re-set for each period t. To avoid that complexity in this illustrative example of the use of the OTS approach, this paper classifies teams according to their location. Teams are classified in five subsets, *North East and Yorkshire*, *North West*, *Midlands* (including here Stoke), *London* and teams that are in locations we can consider as the soccer *Periphery* (including here Swansea, Ipswich, Southampton, Reading etc).

This analysis enables an analysis of whether teams located in different geographic regions perform better (that is, have a lower average rank) than teams in other regions; and whether these relative average rankings change over time.

Table 1 presents the clubs in each region and summarises the results. Based on the average end-of-season rank, London has the most successful teams. London teams have average end-of-season rank of 18th place over the 21 years of the Premier League. The *North West* and *Midlands* are closely matched. Both have an average end-of-season rank of approximately 25th (suggesting the 'average' team would just make Championship playoff positions). *North Eastern* teams finish, on average, 31st in the English league system, and teams from the *Periphery* fare the worst with a notional end-of-season rank for the average team of 35th.

[Table 1 here]

Our methodology enables an appreciation of relative competitive performance temporally, and Figure 3 displays information on changes in regional rankings over time. There is little change in the position of the Midlands (slightly worse) or Peripheral teams (slightly better) over the Premiership period. However, other changes are significant. North Eastern and Yorkshire teams were on average worst performers at the beginning and end of the period, and indeed in addition suffered a longer term decline in relative performance on average – with a monotinic trendline slope of +0.43.

Similarly, while London teams were best performing at both the beginning and end of the period, their performance has improved with respect to other regions; London is not the only region which would not, 'on average', have been relegated in 2013-14, and this is indicative of an evidenced trend. Even more stark is the improvement in performance of North Western teams over the period. From being very much one of the pack in 1992-3, the region is now the only one near to challenging the performance of London.

[Figure 3 Here]

Conclusions

This paper makes a significant contribution by proposing a new tool for measuring league competitiveness which has not been used in sports economics literature but which is directly related to team objectives and which also gives insights into the drivers of team performance. The paper proposes a measure of league competitiveness based on end-of-season league position as a complement to traditional approaches. While our ordinal measure sacrifices a significant amount of

data on the *spread* of performance within leagues, it has benefits relative to traditional win percentages-based measures.

Ordinal measures are more closely aligned with the core team objective – that is, to finish as high in a league as possible - and facilitate easier analysis of relative competitiveness across time, geography and for subsets of teams. Their relative simplicity makes it easy to link league competitiveness and the outcomes of individual teams, enabling a sophisticated criteria analysis, as well as retaining the ability to compare between leagues.

Traditional measures of league competitiveness have developed in a context where, due to franchising and conference-play off structures, ordinal ranking has been impossible or inappropriate. However, where an extended and singular hierarchy exists within a national professional sport, an ordinal measure enables a better understanding of the link between league and team characteristics, and outcome performance.

In addition, such a measure opens up a number of additional analytical avenues, which this paper has only begun to explore the possibilities. These include sophisticated criteria analysis, expanding the approach to other sports leagues, and integrating with team-specific non-competition data (such as on financial power and ownership models for example). A potential important contribution of an ordinal measure is to facilitate analysis of important break points in league hierarchies, such as Champions League qualification, promotion and relegation and their effects on short and long-term team performance and league competitiveness.

References

Bandt, C. (2005). Ordinal time series analysis. *Ecological modelling, 182*(3), 229-238.

Bartimote-Aufflick, K., & Thomson, P. C. (2011). The analysis of ordinal time-series data via a transition (Markov) model. *Journal of Applied Statistics, 38*(9), 1883-1897.

Bowman, R. A., Lambrinos, J., & Ashman, T. (2012). Competitive balance in the eyes of the sports fan: prospective measures using point spreads in the NFL and NBA. *Journal of Sports Economics*, 1527002511430230.

Bridgewater, S., Kahn, L. M., & Goodall, A. H. (2011). Substitution and complementarity between managers and subordinates: Evidence from British football. *Labour Economics, 18*(3), 275-286.

Curran, J., Jennings, I., & Sedgwick, J. (2009). 'Competitive Balance'in the Top Level of English Football, 1948–2008: An Absent Principle and a Forgotten Ideal. *The International Journal of the History of Sport, 26*(11), 1735-1747.

Dobson, S., & Goddard, J. (2011). *The economics of football*: Cambridge University Press.

Fort, R., & Maxcy, J. (2003). "Competitive Balance in Sports Leagues: An Introduction". *Journal of Sports Economics, 4*(2), 154-160. doi: 10.1177/1527002503004002005

Gilmore, S., & Gilson, C. (2007). Finding form: elite sports and the business of change. *Journal of Organizational Change Management, 20*(3), 409-428.

Haskins, M. E. (2013). Oft-overlooked business wisdom highlighted from the youth-soccer sidelines. *Management Decision, 51*(5), 921-937.

Keidel, R. W. (1985). *Game plans: Sports strategies for business*: Beard Books.

Michie, J., & Oughton, C. (2005). Competitive balance in football: An update. *London: the sports Nexus*.

Neale, W. C. (1964). The peculiar economics of professional sports: A contribution to the theory of the firm in sporting competition and in market competition. *The Quarterly Journal of Economics*, 1-14.

Owen, P. D., & King, N. (2014). Competitive Balance Measures in Sports Leagues: The Effects of Variation in Season Length. *Economic Inquiry*.

Owen, P. D., Ryan, M., & Weatherston, C. R. (2007). Measuring competitive balance in professional team sports using the Herfindahl-Hirschman index. *Review of Industrial Organization, 31*(4), 289-302.

Palacios-Huerta, I. (2014). *Beautiful Game Theory: How Soccer Can Help Economics*: Princeton University Press.

Pinch, S., & Henry, N. (1999). Paul Krugman's geographical economics, industrial clustering and the British motor sport industry. *Regional Studies, 33*(9), 815-827.

Rottenberg, S. (1956). The baseball players' labor market. *The Journal of Political Economy*, 242-258.

Ruefli, T. W. (1990). *Ordinal time series analysis: methodology and applications in management strategy and policy*: Quorum Books.

Ruefli, T. W., & Wilson. (1987). Ordinal Time-Series Methodology for Industry and Competitive Analysis. *Management Science, 33*(5), 640-661. doi: DOI 10.1287/mnsc.33.5.640

Szymanski, S. (2001). Income inequality, competitive balance and the attractiveness of team sports: Some evidence and a natural experiment from English soccer. *The Economic Journal, 111*(469), 69-84.

Szymanski, S., & Zimbalist, A. S. (2006). *National pastime: How Americans play baseball and the rest of the world plays soccer*: Brookings Institution Press.

Whitney, J. D. (1988). Winning games versus winning championships: The economics of fan interest and team performance. *Economic Inquiry, 26*(4), 703-724.

Zimbalist, A. S. (2002). Competitive balance in sports leagues: An introduction. *Journal of Sports Economics, 3*, 111-121.

Figure 1: Competitiveness among Premier League Teams 1993-2014 (Teams with >0 Premier League Appearances

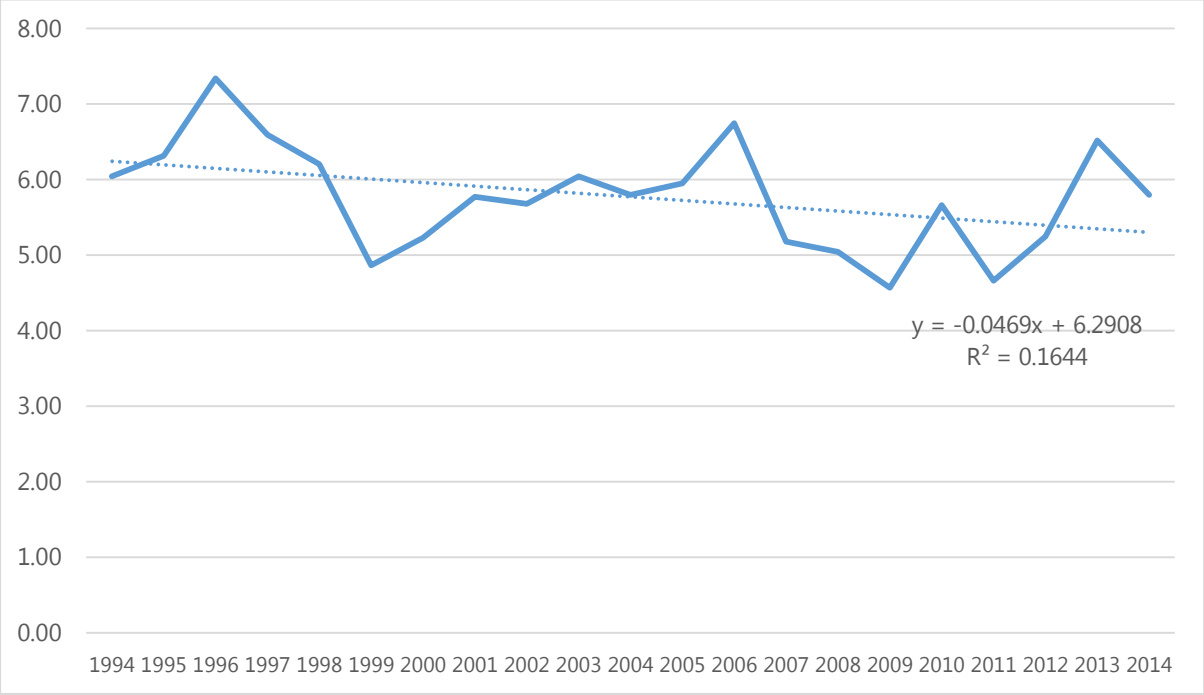

Figure 2: Competitiveness within the English Premier League 199-2014

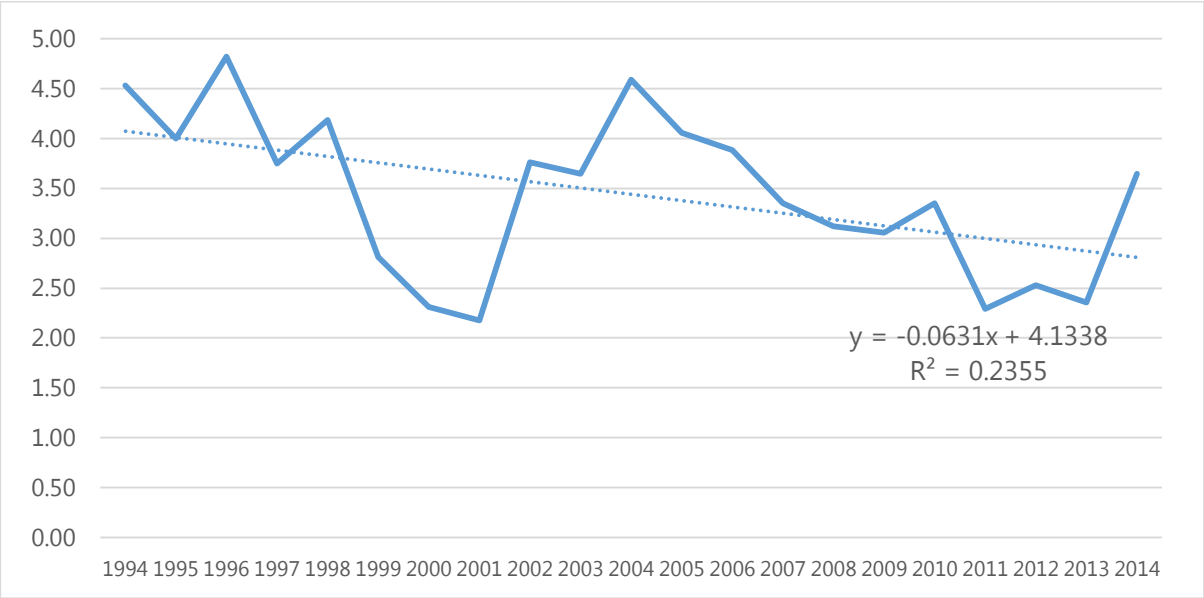

Table 1 - Average Finishing Position 1993-2014 Clubs With >0 Premiership Appearances

Number of Teams	Region	Average End-of Season Position
9	North East & Yorkshire	30.7
8	London	18.4
9	Midlands	25.3
10	North West	24.5
8	Peripheral	34.6

Figure 3: Performance and Geography in the English Soccer Leagues 1993-2014 (Clubs with >0 Premier League appearances)